

ASSISTANT COMMISSIONER FOR PATENTS
BOX PATENT APPLICATION
Washington, D.C. 20231 jc806 U.S. PTO

10-30-00 A
CASE DOCKET NO. CH919990038US1
Date: October 27, 2000

Sir:



10/27/00

under Rule 1.53(f)

Transmitted herewith for filing is the Patent Application/of:

Inventors: Kai Cieliebak and Beat Liver

For: Portfolio Theory Method of Managing Operational Risk with Respect to
Network service-Level Agreements
Enclosed are:

12 (twelve) Sheets of Informal Drawings.

An assignment of the invention to International Business Machines Corporation, Armonk, New York 10504.

A certified copy of a _____ application.

Unsigned Declaration and Power of Attorney is attached to the application.

Associate Power of Attorney.

Information Disclosure Statement with form PTO-1449 with references attached.

The filing fee has been calculated as shown below:

	(Col. 1)	(Col. 2)
FOR:	NO. FILED	NO. EXTRA
BASIC FEE	[REDACTED]	[REDACTED]
TOTAL CLAIMS	29 - 20 =	9
INDEP CLAIMS	8 - 3 =	5
<input checked="" type="checkbox"/> MULTIPLE DEPENDENT CLAIM PRESENTED		

OTHER THAN A SMALL ENTITY	
RATE	FEES
[REDACTED]	\$ 710.00
9 \$ 18 =	\$ 162.00
5 \$ 78 =	\$ 390.00
+ \$ 260=	\$ 260.00
TOTAL \$ 1522.00	

If the difference in Col. 1 is less than zero, enter "0" in Col. 2.

Please charge my Deposit Account No. 09-0468 in the amount of \$ 1522.00.

The Commissioner is hereby authorized to charge payment of the following fees associated with this communication or credit any overpayment to Deposit Account No. 09-0468. A duplicate copy of this sheet is enclosed.

Any additional filing fees required under 37 CFR 1.16.

Any patent application processing fees under 35 CFR 1.17.

Respectfully submitted,

By

David M. Shofri
Registration No.: 39,835
Tel. (914) 945-3247

IBM CORPORATION
INTELLECTUAL PROPERTY LAW DEPT.
P.O. BOX 218
YORKTOWN HEIGHTS, NY 10598

Express Mail Label:EL549238338US
Date of Deposit: October 27, 2000

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Cieliebak et al.

Serial No.: To be assigned

Filed: Herewith

For: PORTFOLIO THEORY METHOD OF MANAGING OPERATIONAL RISK WITH RESPECT TO NETWORK SERVICE-LEVEL AGREEMENTS

Docket No.: CH919990038US1

Group No.:

Examiner:

Commissioner of Patent and Trademarks
Washington, DC 20231



EXPRESS MAIL CERTIFICATE

Express Mail Label Number EL549238338US
Date of Deposit October 27 2000

I hereby certify that attached paper or fee:

Patent Application Transmittal Fee Sheet under Rule 1.53(f) -Original and copy
Patent Application
Declaration and Power of Attorney for Patent Application **unsigned**
12 (twelve) Sheets Informal Drawings
Return Post Card

is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and addressed to the Commissioner of Patents and Trademarks, Washington, DC 20231.

Margaret McCabe

(Name)

Signature

Note: Each paper must have its own certificate and the "Express Mail" label number as a part thereof or attached thereto. When, as here, the certification is presented on a separate sheet, that sheet must (1) be signed and (2) fully identify and be securely attached to the paper or fee it accompanies. Identification should include the serial number and filing date of the application as well as the type of paper being filed, e.g. complete application, specification and drawings, responses to rejection or refusal, notice of appeal, etc. If the serial number of the application is not known, the identification should include at least the name of the inventor(s) and the title of the invention.

Note: The label number need not be placed on each page. It should, however, be placed on the first page of each separate document, such as, a new application, amendment, assignment, and transmittal letter for a fee, along with the certificate of mailing by "Express Mail". Although the label number may be on checks, such a practice is not required. In order not to deface formal drawings it is suggested that the label number be placed on the back of each formal drawing or the drawings be accompanied by a set of informal drawings on which the label number is placed.

1 PORTFOLIO THEORY METHOD OF MANAGING OPERATIONAL
2 RISK WITH RESPECT TO NETWORK SERVICE-LEVEL
3 AGREEMENTS

4 CROSS REFERENCE TO RELATED APPLICATION

5 This application claims priority to co-pending U.S.
6 provisional application no. 60/162,383 filed October 28, 1999.

7 Field of the Invention

8 This invention relates to a method of managing risk, and
9 more particularly, to a method of managing operational risk and return
10 with respect to network service-level agreements ("SLA"s).

11 Background of the Invention

12 In order to ensure economical network operations, providers
13 are concerned with the following trade-off: on the one hand, better
14 Quality of Service corresponds to higher price, thereby increasing
15 revenue. On the other hand, if the provider guarantees higher Quality of
16 Service and is not willing to run a higher risk, he can only accept less
17 traffic, thereby decreasing revenue. In order to properly evaluate this
18 trade off, the provider attempts to manage operational risk associated
19 with non-complying network service-level agreements.

20 In the prior art, operators employ simple traffic engineering
21 to meet the QoS as specified in the SLAs. For example, sensitivity
22 analysis is carried out to determine the likelihood of violating SLAs.

1 Therefore, what is needed is a systematic method to evaluate
2 risk and return with respect to network service-level agreements that can
3 be implemented on a computer in order to provide near real time
4 assessments of performance, thus providing more accurate risk
5 assessments and less uncertainty.

Summary of the Invention

7 It is therefore an object of the invention to provide a method
8 of managing operational risk and return with respect to a portfolio of
9 classes of service-level agreements (“SLA”s). The method executes the
10 following steps: (1) calculating an efficient frontier that identifies
11 efficient portfolios of SLAs using inputs such as characteristics of the
12 production infrastructure, traffic and QoS characteristics and the price of
13 each class of SLAs; (2) optionally, calculating a baseline efficient
14 frontier using inputs such as market pricing and break-even (zero-profit)
15 pricing; (3) determining the performance of the current portfolio of
16 SLAs using a portfolio evaluator means and inputs which characterize
17 the current portfolio; (4) evaluating performance by comparing the
18 current portfolio and the efficient portfolios with the desired level of risk
19 and return; and, if desired, implementing corrective action based on any
20 desired risk and return.

Brief Description of the Drawings

22 The above brief description, as well as further objects,

1 features and advantages of the present invention will be more fully
2 appreciated by reference to the following detailed description of the
3 presently preferred but nonetheless illustrative embodiments in
4 accordance with the present invention when taken in conjunction with
5 the accompanying drawings.

6 FIG. 1 is a flow diagram of the method of the invention.

7 FIG. 2 is a schematic diagram of a computer device on
8 which the invention operates.

9 FIG. 3 is a diagram of a network on which the invention
10 may be implemented.

11 FIG. 4 is a detailed flow diagram of the method of the
12 invention.

13 FIG. 5 is a schematic view showing portfolio theory applied
14 to network operations.

15 FIG. 6 is a flow chart illustrating a Portfolio Evaluator of
16 the invention.

17 FIG. 7 is a graph of risk vs. return showing the efficient

1 frontier.

2 FIG. 8 is a schematic diagram of an SLA.

3 FIG. 9 is a graph showing examples of extremal points.

4 FIG. 10 is a graph of a polyhedron of constant return.

5 FIG. 11 is a schematic view of an example ring network.

6 FIG. 12 is a graph of the normalized traffic distribution X.

7 FIG. 13 is a zero-profit price curve.

8 FIG. 14 is a graph of the risk and return of portfolios.

9 Detailed Description of the Invention

10 Glossary of Terms and Symbols

11 βC financial penalty per capacity unit.

12 C capacity C (of stem, network, link and so on)

13 \mathbf{D} vector of Quality of Service classes, in case of delay

14 D_i Quality of Service offered by class i in case of delay

15 $e(y)$ return of a portfolio y

1 \underline{L} vector of Quality of Service classes, in case of loss ratio
 2 L_i vector of Quality of Service classes in case of loss ratio
 3 $p_C C$ constant term reflecting the marginal cost of providing the
 4 network.
 5 P denotes a portfolio
 6 \underline{p} price vector
 7 p_C unit price for capacity C
 8 p_i price of contract of type i (expected revenue)
 9 \underline{q} contracted Quality of Service of the contracts of the portfolio;
 10 \underline{q}' (expected) actual Quality of Service of a network
 11 $r(\underline{y})$ risk of a portfolio \underline{y}
 12 $r_{QoS}(\underline{y})$ Quality of Service risk, i.e., risk expressed in terms of QoS units
 13 $r_s(\underline{y})$ financial risk, i.e., risk expressed in terms of monetary units
 14 \mathbf{R} rational numbers
 15 \mathbf{R}_+ rational numbers that are greater than 0
 16 \mathbf{R}^n n-dimensional space, where each dimension is of \mathbf{R}
 17 \underline{y} a portfolio, i.e., $\underline{y} = \langle y_1, \dots, y_i, \dots, y_n \rangle$
 18 y_i amount of contracts (SLAs) of type i

19 Referring now to FIG. 1, which is a flow diagram of the
 20 invention, the invention provides a method 10 and a system 20 that
 21 applies the principals set forth in detail in provisional application no.
 22 60/162,383, hereby incorporated by reference. The method 10 manages
 23 operational risk and return with respect to a portfolio of classes of

1 computer resource or service-level agreements (“SLA”s) by executing
2 the following steps. In a first step 12, the method 10 calculates an
3 efficient frontier 110 that identifies efficient portfolios of SLAs using
4 inputs such as characteristics of the production infrastructure 138, traffic
5 and QoS characteristics and the price of each class of SLAs. In a second
6 step 14, the method 10, optionally, calculating a baseline efficient
7 frontier 110 using inputs such as market pricing and zero-profit pricing.
8 In a third step 16, the method 10 determines the performance of the
9 current portfolio of SLAs using a portfolio evaluator 144 and inputs that
10 characterize the current portfolio. In a fourth step 18, the method 10
11 evaluates performance by comparing the current portfolio and the
12 efficient portfolios with the desired level of risk and return; and, if
13 desired, implements corrective action based on any desired risk and
14 return.

15 Referring now to FIG. 2, which is a schematic diagram of a
16 typical system 20 for practicing the various embodiments of the present
17 invention, the method 10 is encoded on a computer-readable medium
18 and operates on a computer system 20 and/or between the computer
19 system and a server 25 or 54 (shown in FIG. 3) on an intranet or the
20 Internet. Such a computer system 20 typically includes a computer 22, a
21 display device 24, an input device 26 such as a keyboard, a primary
22 storage device 30 and a secondary storage device 32. After loading of
23 software encoded with the method 10 of the invention or after accessing
24 the server 25 or 54 through a browser such as Internet Explore 5.0, as the
25 case may be, the display device 24 displays a graphical user interface

1 ("GUI") 34 for facilitating the display of text and graphics associated
2 with the method to the user. Display devices 24 include printers and
3 computer display screens such as a CRT, LED displays, LCDs, flat
4 screens, screen phones, and projectors. Input devices 26 are numerous
5 and include keyboards and pointing devices such as a mouse 27 having a
6 left mouse button 28 and a right mouse button 29, a trackball, lightpens,
7 thumbwheels, digitizing tablets, microphones using voice recognition
8 software, and touch screens and pads.

9 The GUI 34 provides input fields for data input and control
10 of the method 10, as well as an output window for statistical displays of
11 information, which facilitates management of the network. The method
12 10 accesses a database in primary storage 30, the database including
13 information associated with each SLA, organized in a data structure
14 including the class i of the SLA, the terms 126 of each SLA, such terms
15 including the offered capacity 122, the Quality of Service guarantees 124
16 with respect to delay, loss, and availability, a price 126, a penalty 130, a
17 duration 132, and, optionally, relative compliance guarantee(s) 86a
18 (shown in FIG. 8).

19 The computer 22 includes a CPU 36 as well as other
20 components with which all who are skilled in the art are familiar. For a
21 detailed discussion of these components and their interaction, see U.S.
22 Pat. No. 5,787,254, the content of which is incorporated by reference.
23 The secondary storage 32 supports the method 10, preferably
24 HTTP-compliant, as well as a number of Internet access tools. The
25 CPU 36 fetches computer instructions from primary storage 30 through

1 an interface 40 such as an input/output subsystem connected to a bus 42.
2 The computer 22 can be, but is not limited to, an "IBM APTIVA"
3 computer, a product of International Business Machines Corporation of
4 Armonk, New York, or any computer compatible with the IBM PC
5 computer systems based on the X86 or Pentium(TM) series processor of
6 Intel Corporation or compatible processors, or any other suitable
7 computer. The CPU 36 utilizes an operating system that, depending on
8 the hardware used, may be DOS, "WINDOWS 3.X", "WINDOWS
9 XXXX", "NT", "OS/X", "AIX", "LINUX", or any other suitable
10 operating system. The CPU 36 executes these fetched computer
11 instructions. Executing these instructions enables the CPU 36 to retrieve
12 data or write data to the primary storage 30, display information, such as
13 the statistical displays of the method 10, on one or more display devices
14 24, receive command signals from one or more input devices 26, or
15 transfer data to secondary storage 32 or even other computer systems
16 which collectively form a computer network 25 (shown in FIG. 3).
17 Those skilled in the art understand that primary storage 30 and
18 secondary storage 32 can include any type of computer storage including
19 RAM, ROM, application specific integrated circuits ("ASIC") and
20 storage devices that include magnetic and optical storage media such as
21 a CD-ROM.

22 Where the method 10 operates on a stand-alone computer
23 22, the primary storage 30 stores a number of items including the
24 method 10 and a runtime environment 46. The runtime environment 46
25 typically is an operating system that manages computer resources, such

1 as memory, disk or processor time, required for the method of the
2 invention to run. The runtime environment 46 may also be a message
3 passing system, a microkernel, dynamic loadable linkable module(s), or
4 any other system that manages computer resources.

5 Now referring to FIG. 4, in which a more detailed flow
6 diagram of the method is shown, the method 10 includes the following
7 steps. In a first step 60, the method gathers inputs from the provider
8 including characteristics of the production infrastructure, the QoS
9 characteristics and price of each possible and reasonable class of SLA.
10 In a second step 62, the method 10 calculates an efficient frontier 110
11 (shown in FIG. 7) that identifies efficient portfolios of SLAs.
12 Optionally, the method 10 substitutes the actual pricing of SLAs with
13 baseline pricing such as market pricing or break-even pricing, in order
14 for the operator to obtain insights regarding the effects of price changes
15 on his risk and return, with respect to the market. In a third step 64,
16 which may run concurrently with the first and second steps 60 and 62,
17 the method 10 gathers inputs characterizing the current portfolio of
18 SLAs and the desired risk and return. In a fourth step 66, which may run
19 concurrently with the first and second steps 60 and 62, the method 10
20 computes the risk and return of the current portfolio using a portfolio
21 evaluator 144 (shown in FIG. 6). In a fifth step 70, the method 10
22 calculates the difference between the optimal portfolio identified by the
23 efficient frontier 110 and the current portfolio. In a sixth step 72, the
24 difference is evaluated. If actual risk and return matches the desired

1 levels, then an acceptable portfolio 74 has been attained and the method
2 waits a period of time ΔT (depicted in the figure by box 76), before
3 restarting the method. Otherwise, in a seventh step 80, if actual risk is
4 higher than desired risk or if actual return is lower than desired return,
5 the method 10 takes corrective action. Corrective action can include
6 adjusting marketing strategy 82, changing the degree of multiplexing,
7 84, defining relative compliance guarantees and running packets through
8 a service discipline which allows transmission on the basis of priority (as
9 defined by the guarantees specified in the SLAs), 86, changing prices ,
10 90, trading different classes of SLAs, 92, and/or reducing the costs of the
11 production infrastructure 94. In a seventh step 96, after an adjustment
12 due to the selected corrective action is made to the production
13 infrastructure, the method 10 takes new inputs, and, with the exception
14 of the corrective action of trading SLAs, 92, the method is re-executed,
15 by calculating a new efficient frontier 110 which is compared with actual
16 performance, calculated by the portfolio evaluator 144, given the new
17 parameters.

18 Portfolio Theory and Service-Level Agreements

19 In calculating the efficient frontier 110, the method 10
20 applies the principles of classical Portfolio theory — to be precise, the
21 pre-CAPM (Capital Asset Pricing Model) version of portfolio theory,
22 which was initially developed by H. Markowitz, W. Sharp and others for
23 portfolios of classes of financial assets (shares, bonds, etc.), to provide a

1 framework in which to describe this trade-off between risk and return for
2 portfolios of classes of SLAs. In the classical application of portfolio
3 theory, it is assumed that there are finitely many assets i .

4 Each SLA in the portfolio specifies a peak rate (e.g., bits per
5 second) and a Quality of Service guarantee (e.g., loss rate). Associated
6 with each portfolio is its return (relative profit) and its risk of violating
7 any of the SLAs. This risk will be referred to as non-compliance risk (the
8 risk that any of the Quality of Service guarantees of the sold SLAs is
9 violated). In contrast to return, risk generally cannot be quantified in
10 monetary terms directly. Quantifying risk in monetary terms requires two
11 steps:

12 1. Risk is measured in quantities specific to the asset.
13 2. The measured risk levels have to be valued in terms of the
14 contract value (e.g., money-back guarantee) specified in the contracts.

15 In order to separate these two steps and apply different
16 valuation methods, risk and return are treated as independent parameters
17 associated with portfolios.

18 Assuming that the set of attainable portfolios is all
19 nonnegative real numbers \mathbf{R}_+ up to the number n of available assets
20 (which is finite), each portfolio may be associated with two quantities:
21 the (expected) return and the risk. A portfolio is called efficient if it
22 maximizes return at a given risk, or equivalently, minimizes risk at a

1 given return. The efficient frontier 110 is the image in the risk-return
2 space of the set of efficient portfolios.

3 Treating the contracts, e.g., SLAs, of a service provider as a
4 portfolio makes it possible to develop decision-support tools for
5 determining the Quality of Service classes to be offered and for
6 managing the noncompliance risk — an operational risk resulting from
7 multiplexing (and other behaviors of a production infrastructure in
8 operation). One such management strategy, corrective action 86, defines
9 a new contractual parameter called relative compliance guarantees,
10 which will be discussed in more detail later, along with a discussion of
11 the trading of risks.

12 Referring now to FIG. 5, a network service provider sells
13 connectivity over a network 100. The provider offers n classes of
14 service level agreements $SLA_1 \dots SLA_n$. Each SLA specifies a
15 connection, contract duration, traffic descriptors (peak rate, average rate,
16 burst size, etc.) and Quality of Service guarantees (loss rate, delay, jitter,
17 etc.). An SLA is normalized to a peak rate (or average rate) of 1 bit/s.
18 Then the provider's portfolio of SLAs is given by a vector $\underline{y} \in \mathbf{R}_+^n$
19 whose component y_i is the number of contracts of class i .

20 Associated with each portfolio are two quantities: the
21 (expected) return $e(\underline{y})$ and the risk $r(\underline{y})$. The return (or profit) of a
22 portfolio \underline{y} equals $e(\underline{y}) = \sum p_i y_i - p_c C$ where p_i , C and p_c denote the unit
23 price of class i , the capacity of the network and the unit price of network
24 capacity respectively. Note that $p_c C$ is a constant term reflecting the

1 marginal cost of providing the network. The unit price p_c depends on C
2 as networks 100 exhibit economies of scale in general. In the case of a
3 single link, C is the link capacity.

4 A portfolio γ entails a risk of noncompliance $r(\gamma)$ for the
5 provider that depends on the traffic statistics (i.e., the traffic that is
6 actually sent by consumers under their SLAs within the specified traffic
7 descriptor), as well as on the network topology and capacities. There are
8 different risk measures conceivable (as discussed below).

9 In order to help structure his portfolio γ of SLAs, the
10 provider must consider as inputs such factors as traffic statistics 102,
11 market information 104, and the structure and behavior of the network
12 100. Then, by evaluating risk and return 106, he may determine the
13 efficient frontier 110 (discussed in detail in connection with FIG. 7).

14 The set of feasible portfolios 112 (shown in FIG. 7) and the
15 prices p_i will be determined by the market demand. Once a network
16 service provider has determined the appropriate risk measure, which may
17 be any risk measure, and has derived a way to compute it, he can think
18 about his operations in the terms of portfolio theory. Doing so enables
19 the provider to (1) decide how many and which types of SLAs to offer
20 (described above); (2) evaluate the efficiency of the current portfolio; (3)
21 compute the efficient frontier 110; (4) quantify risk and return 106 of the
22 current portfolio; (5) derive strategies to move towards a more efficient
23 portfolio, and (6) determine base-line portfolios for (cost-based)
24 zero-profit prices.

1 Evaluate the efficiency of the current portfolio .

2 In order to obtain the performance characteristics of the
3 existing production infrastructure 138, for comparison with the efficient
4 portfolio 110 (i.e., the fourth step 18 of method 10, shown in FIG. 1), a
5 Portfolio Evaluator 144 is provided. In addition to portfolio details and
6 the production infrastructure (characterized by the vector \underline{i} , which is
7 fixed here and hence not further discussed), the Portfolio Evaluator 144,
8 shown in FIG. 6, takes a Boolean variable “ S ”, as input to select between
9 risk measure 136a, “ $r_s(y)$ ”, and risk measure 136b, “ $r_{QoS}(y)$ ” (i.e., the
10 provider decides whether he wishes to evaluate the risk of a penalty or
11 the risk of violating a Quality of Service requirement) . The Portfolio
12 Evaluator 144 carries out the following steps:

13 (1) A Performance Evaluator 146 is invoked to determine the
14 (expected) actual Quality of Service 150, “ q ”. The Performance
15 Evaluator 146 is a formula (if an analytical performance model
16 exists) or a simulator. Further, the actual details of the
17 infrastructure 138 may be used for determining performance.

18 (2) The portfolio risk 136, “ $r(y)$ ”, is computed based on actual Quality
19 of Service 150, q' , and the contracted Quality of Service 152, q , of
20 the contracts of the portfolio using a particular *risk measure* 136a
21 or 136b.

22 (3) The return 134, “ $e(y)$ ”, is computed according to the formula 154,

1 $\sum p_i y_i - p_C C$, for capacity 140, “C”, and capacity unit price 142,
2 “ p_C ”, where capacity 140 is an input in both the Performance
3 Evaluator 146 and is characteristic of the production infrastructure
4 138.

5 Risk measures 136 include the *risk of noncompliance* (i.e.,
6 the risk of not being able to satisfy all the Quality of Service guarantees
7 of the sold contracts expressed as the probability that some SLAs are
8 violated) and the expected excess quality for class i (i.e., the expected
9 value of the difference between the delivered and contracted Quality of
10 Service for class class i). Clearly, there are many such risk measures
11 136, which may be conceived. Alternatively, if the SLAs contain explicit
12 penalties for noncompliance, risk can be measured as the expected
13 penalty due to contract violations. Which risk measure 136 is more
14 appropriate depends on the business implications of noncompliance: for
15 large customers pursuing long-term relationships with the provider, the
16 provider will strive to comply with all contracts, so he wishes to keep the
17 probability of noncompliance at a low level. On the other hand, for
18 small consumers frequently changing providers, the provider may
19 deliberately risk contract violations, incorporating expected penalty as a
20 cost in his profit function. Note that the current portfolio can be
21 evaluated using the measured performance of the infrastructure, i.e., the
22 Performance Evaluator 146 is a database of performance data. Such data
23 is typically available from performance management studies or reports.

1 Determining risk measure 136a, $r_S(y)$, based on specified
2 penalties is just one method to value the risk measure 136b, $r_{QoS}(y)$, in
3 financial terms, called a *valuation method*. Alternative methods are
4 conceivable including the use of quantified user satisfaction based on,
5 for instance, surveys and experiments. This satisfaction might depend on
6 the market segment (e.g., business and private customers), so that it
7 would be necessary to assign different values to each such group of
8 contracts. A second alternative is given below.

9 In step 80 of method 10, a provider finding out that his risk
10 136b, $r_{QoS}(y)$, is not zero — he sometimes violates some SLAs — takes
11 corrective action. In corrective action 92, he may re-engineer his
12 infrastructure including increasing the capacity C or accept the risk and
13 pay penalties, if such are specified, or accept unsatisfied customers.
14 Contracts with particularly high Quality of Service guarantees require
15 more resources to guarantee them. However, these high capacity
16 requirements are offset when portfolio mixes such high Quality of
17 Service requirement SLAs with contracts that offer only a low Quality of
18 Service (e.g., a high loss rate) or a low probability of compliance. This
19 leads either to a higher return, lower risk or lower price (or a
20 combination therefore).

21 Referring now to FIG. 4, in corrective action 86, wherein
22 relative compliance guarantees are used, the method 10 of the invention
23 implements a service discipline 86b which allows the degradation the
24 Quality of Service of a communication flow according the Quality of

1 Service specified in the corresponding SLA. The service discipline 86b
2 is carried out by the network (assuming that it is possible to program the
3 network for this purpose) Thus, the provider offers SLAs that guarantee
4 relative compliance 86a. These relative compliance guarantees 86a are
5 specified in terms of a noncompliance risk measure. In practice, a
6 premium is charged for the higher compliance probability. Compliance is
7 hence a product differentiator — a measure which may become as
8 important as network reliability. Note that a portfolio y containing SLAs
9 with relative compliance guarantees 86a can be evaluated with the same
10 approach to evaluate whether these relative compliance guarantees are
11 met. Therefore, the method 10 provides this new contractual parameter,
12 *relative compliance guarantees* 86a. The contractual parameter is
13 calculated in step 70 of the method 10, in which the difference between
14 the actual and the desired risk is equated to the relative compliance
15 guarantee, which is added as a SLA contractual parameter, to define a
16 lower service level. Making the risk explicit enables new valuation
17 methods that, in particular, take advantage of the willingness of
18 consumers to pay a certain amount for a given risk level.

19 Compute the efficient frontier

20 Referring now to FIG. 7, portfolio theory is concerned with
21 the computation and properties of the efficient frontier 110. Once the
22 efficient frontier 110 has been determined, it is a business decision to
23 select a portfolio on the efficient frontier, depending on the tolerable

1 level of risk or the target return.

2 In order for the provider to gain an insight into where his
3 current portfolio stands with respect to an efficient portfolio that
4 maximizes profit for a given risk, steps 60 and 62 of the method 10 apply
5 the principals of Portfolio Theory to calculate the efficient frontier 110.
6 FIG. 7 shows the return-risk space with the attainable portfolios and the
7 set of efficient portfolios, i.e., the efficient frontier 110. The efficient
8 frontier 110 is defined by a closed-form formula, which is only possible
9 in special cases. Assessing the efficiency of the current portfolio P^*
10 requires the computation of the efficient frontier 110. The example
11 shown in the figure consists of three segments: two of them result from
12 pairs of adjacent extremal points (shown in FIG. 9 , identified in a closer
13 analysis of the quasi-linearity of the risk function in Portfolio Theory),
14 and the third consists of portfolios of a single Quality of Service class.

15 It is assumed that return 134 is a linear function (as defined
16 above), equal to the summation of the product of each vector describing
17 the portfolio multiplied by nonnegative coefficients of a price vector
18 associated with each vector describing the portfolio, from which
19 marginal cost (a constant) is subtracted.

20 Risk measures 136 can be characterized as convex and
21 quasi-linear risk functions. A function is called convex if all sublevel
22 sets are strictly convex, which yields the following implication, *Lemma*
23 *1*: If the risk measure is a convex risk function, then for every price
24 vector and risk level, there exists a unique portfolio that maximizes

1 return at a given risk level. The function that describes the efficient
2 portfolios is continuous in both the certain risk level and in the price
3 vector. The amount of the asset in the unique portfolio is zero whenever
4 the price vector associated with that asset is also zero.

5 A risk function r is called quasi-linear if it depends only on
6 the two quantities, the summation of y_i , the vector description of an SLA
7 in a portfolio and the summation of the product of loss rate L_i for a
8 particular SLA i and y_i , for some vector $\underline{L} = (L_1, \dots, L_n) \in \mathbf{R}_n^+$ which
9 characterizes the quality of each SLA (where the lower loss ratio L_i
10 corresponds to better quality). Note that instead of $\sum y_i$, any linear
11 function $\sum M_i y_i$ with positive coefficients $M_i > 0$, could have been used
12 because the transformation $y_i \rightarrow M_i y_i$, $p_i \rightarrow p_i/M_i$ shows that this is
13 equivalent to the case where all $M_i = 1$. If risk is expressed in terms of a
14 special function of c , the inverse of the vector of an asset and the loss
15 ratio, then the condition that the partial derivative of the special function
16 with respect to c and the partial derivative with respect to the loss ratio
17 are less than zero ensures that the risk increases with the aggregate of
18 assets as well as with the quantity.

19 If n , the number of classes of SLAs, is less than 2, a
20 quasi-linear risk function cannot be convex in the same sense as
21 described above. The fact that the special function is convex provides
22 the best proxy of convexity for a quasi-linear risk function.

23 Referring now to FIG. 9, quasi-linearity has the following
24 consequence: *Lemma 2:* for a quasi-linear risk function, then (i) the

1 efficient frontier 110 is generated by portfolios consisting of one or two
2 classes of SLAs ; (ii) a portfolio consisting of one SLA i is efficient
3 only if (L_i, p_i) , the loss ratio for the SLA and the unit price for that SLA,
4 constitutes an extremal point on the graph of the price $p(L)$ vs. L , loss
5 ratio shown in FIG. 9 , i.e., it lies on the boundary of the curve
6 representing the convex hull in the graph of price vs. loss ratio (therefore,
7 a portfolio of two SLAs, i , and j , is efficient only if (L_i, p_i) and (L_j, p_j) are
8 adjacent extremal points); (iii), supposing that the special function is
9 convex, then there exists a function that assigns to every price vector and
10 risk level greater than or equal to zero, an efficient portfolio of a certain
11 risk consisting of one or two SLAs; and (iv), for a number of SLAs
12 exceeding 2, a function as in “(iii)” cannot be continuous everywhere.

13 Model 1: Loss

14 Assuming that the Quality of Service is described by a
15 single parameter, the loss ratio L , defined as the proportion of lost bits to
16 sent bits in a given time interval of duration T , the relations developed
17 above can be illustrated with a real world example, Model 1, in which
18 the network consists of a single link of capacity C . This is useful due to
19 the fact that single links are important as access lines (e.g., an xDSL line
20 connecting a customer site with a central office) and hot spots, and will
21 be discussed in further detail below. Further, the method 10 assumes that
22 the network employs a proportional scheduling service discipline 86a
23 which ensures that whenever the aggregate condition, defined by the

1 total lost traffic being less than or equal to the summation of the product
2 of the loss ratio L_j and the random variable, X_j , denoting the traffic sent
3 by customers of class j , holds, the lost traffic for each contract does not
4 exceed the specified loss ratio. Then, assuming further that there exists a
5 random variable Y such that $\sum X_i \sim (\sum y_i)Y$ and $\sum L_i X_i \sim (\sum L_i y_i)Y$, where \sim
6 denotes equality in distribution, then the risk function is quasi-linear
7 (depending only on $\sum y_i$ and $\sum L_i y_i$). Therefore, the conclusions (i) and
8 (ii) of *Lemma 2* hold, and one can conclude that the efficient frontier 110
9 is generated by portfolios consisting of at most two Quality of Service
10 classes L_i , L_j corresponding to adjacent extremal points on the price
11 curve.

12 This is consistent with the findings for simple networks of
13 Kai Cieliebak and Beat Liver, in their provisional application in which
14 it was shown that the efficient frontier 110 is generated by portfolios
15 consisting of at most two Quality of Service classes, L_i , L_j , corresponding
16 to adjacent (i.e., a line segment joining them is contained in the boundary
17 of S) extremal points on the price curve of FIG. 9.

18 In case a network has conceptually a common queue of
19 packets (with respect to the considered Quality of Service parameter), a
20 proportional scheduling policy (with the above-described property) exists
21 and hence Lemma 2 holds. Many broadcast network protocols have this
22 property, so that someone skilled in the art can develop the required
23 proportional scheduling policy. For example, the implementation of this
24 policy for the CSMA/CD (Carrier Sense Multiple Access/Collision

1 Detection) is described as follows. First, each network node has to carry
2 out admission control. Second, a network node uses the standard
3 retransmits protocol for dealing with collisions if the lost traffic for class
4 i exceeds the contracted loss ratio multiplied by the traffic sent by class i
5 (i.e., $Z_i > \sum_j y_j$). Otherwise, packets are not retransmitted.

6 For non-broadcast networks, routing must be taken into
7 account. The results for a single link apply only for special cases in
8 which a network can be treated as a set of independent links. One way
9 this can occur is if multiplexing among different flows is prevented.
10 Another possibility is a highly symmetric topology that makes the
11 network equivalent to independent Links as shown in FIG. 11. A ring
12 network 160 consisting of four nodes 160a, 160b, 160c, and 160d, four
13 links 162a, 162b, 162c, and 162d and two flows 164a and 164b between
14 nodes 160a and 160c, and between 160b and 160d. Flow 164a is equally
15 distributed over the two possible paths for Flow 164b, and vice versa.
16 So, for multiplexing purposes, this network 160 is equivalent to a single
17 link shared by the two flows 164a and 164b. In this figure, a dotted line
18 represents aggregate flows X^{ij} . In both paths, all links have the same
19 capacity c_l . The capacity of the ring network 160 depends on $X^{4,8}$: If $X^{4,8}$
20 varies between 0 and c_l and it is routed clock-wise, the capacity available
21 to $X^{1,3}$ and $X^{5,3}$ varies between c_l and 0. Consequently, Z depends on the
22 traffic situation.

23 Model 1 can be applied to real world networks. Networks
24 fall into two broad categories: broadcast networks (e.g., Ethernet and

1 token ring) and networks using point-to-point connections. Some
2 broadcast networks have conceptually a common queue of packets, i.e.,
3 the shared medium may be treated like a single link. For such networks,
4 the equations for noncompliance risk with loss guarantees and expected
5 penalty for loss, given below, apply. In fact, there exists a large
6 number of broadcast networks that can be modeled as a single link.
7 These include Carrier Sense Multiple Access (CSMA), CSMA/CD
8 (Collision Detection) – better known as Ethernet, token buses and rings,
9 wireless networks, and satellite up-links.

10 Quantify risk and return of the current portfolio

11 In the fourth step 66 of method 10, formulas for risk
12 measures are called for. Two specific formulas for quasi-linear risk
13 measures may now be provided. First, the following definitions are
14 made: $y = \sum_i y_i$; $c = C/y$; $L = (\sum_i L_{iy_i})y$, and the random distributions are
15 written as $\mathbf{Z} = (\mathbf{X} - C)^+ \sim (y\mathbf{Y} - C)^+ = y(\mathbf{Y} - c)^+$, $\sum_i Z_i \sim (\sum_i L_{iy_i})\mathbf{Y} = LY$.

16 The probability of noncompliance with loss guarantees equals $PNL(c, L)$
17 $= P[Z > \sum_i X_i] = P[(Y - c)^+ - LY > 0]. \quad (1)$

18 This measure 136 defines the portfolio risk that is the
19 probability that some SLA of the portfolio is violated. Here the pair of
20 variables $(y, \sum_i y_i)$ has been replaced with the equivalent pair (c, L) . The
21 probability of noncompliance can be computed from this formula once

1 the distribution of Y is known (e.g., from historical data).
2 Making the reasonable assumption that the aggregate
3 penalty for noncompliance is proportional to the lost traffic in excess of
4 the SLAs, $Z - \sum_i X_i$, then the expected penalty for loss equals:

5 $EPL(c, L) = (\beta C)E[Z - \sum_i X_i],$ (2)
6 for some constant $\beta > 0$, so that (βC) denotes the penalty per capacity
7 unit.

8 Model 2: Delay

9 In this section, a second model, Model 2, is described that is
10 complementary to the previous one, based on the following two basic
11 assumptions, namely, (1) a single link and (2) the Quality of Service is
12 described by a single parameter, the delay D . Assume that the link serves
13 customers of guaranteed delays $D_1 < \dots < D_n$. As in the preceding
14 sections, the service discipline is activated which customers of class i
15 have strict priority over customers of class $j > i$ (head-of-line), but
16 service in progress is not interrupted (i.e., non-preemptive).

17 In contrast to the preceding sections, where a general
18 scaling assumption was sufficient, here a specific traffic distribution
19 must be assumed: customers of class i arrive at Poisson rate λ_i , and the
20 arrival processes are independent of each other. Further, service times
21 are identically distributed and they are independent of each other and of

1 the arrival processes (M/G/1 queuing system).

2 Under the assumptions that the network consists of one link
3 of capacity C , the Quality of Service is described by a single parameter
4 (the delay D) and the assumption in the above paragraph, the expected
5 penalty for delay, $EPD(c,D)$, is a quasi-linear risk function that is
6 convex. Therefore, conclusions (i) and (iii) of *Lemma 2* hold: The
7 efficient frontier 110 is generated by portfolios consisting of at most two
8 Quality of Service classes D_i, D_j corresponding to adjacent extremal
9 points on the price curve. Moreover, there exists a function $v^\rho(p)$ that
10 assigns to every risk ρ and price vector p a portfolio of at most two
11 Quality of Service classes, which is continuous except at price vectors
12 where the set of extremal points changes.

13 The expected penalty for delay, EPD is computed over a
14 time interval from the formula: $EPD(c, L) = \beta \sum \{(\lambda_i / \mu_i)(E[W_i] - D_i)\} =$
15 $\beta \{ (\alpha/(c-1)) - (D/c) \}$, where β is a constant > 0 , $c = 1 / \sum \{\lambda_i / \mu_i\}$, $D = c \sum$
16 $\{(\lambda_i / \mu_i)D_i\}$, and $E[W_i]$ denotes the expected waiting time (i.e., delay) for
17 class i . Assuming that class i traffic arrives at Poisson rate λ_i and the
18 arrival process are independent of each other; service times,
19 characterized by service rate μ_i of class i , are independently distributed,
20 and they are independent of each other and of the arrival processes —
21 i.e., an M/G/1 queuing system is assumed. Assuming that the service
22 times for customers of all classes are distributed as a random variable Y
23 of mean μ then $\alpha = (1 + \{Var[Y]/\mu^2\})^2/2$, where $Var[Y]$ denotes the
24 variance of random variable Y . Note that noncompliance is defined here

1 in terms of a penalty for exceeding D_i and a premium for remaining
2 under D_i .

3 Determine base-line portfolios for (cost-based) zero-profit prices

4 In step 62 of method 10, determining base-line scenarios, is
5 useful to provide insights in the economics of a network's operation. The
6 method 10 optionally calculates a base-line efficient frontier (or
7 portfolio), assuming that there exists sufficient demand for all
8 considered Quality of Service classes. This means that \mathbb{R}^n_+ defines the
9 set of attainable portfolios. A provider would most likely wish to
10 determine the base-line efficient frontier first. Then, he can investigate
11 which of these portfolios are probably attainable and compare the
12 base-line prices against markets prices (e.g., to determine which Quality
13 of Service classes to offer).

14 For base lining, the prices p_i can be defined as zero-profit
15 prices at the risk level $EPL(c, L)=0$ — so that profit equals costs — by
16 setting prices proportional to the resource consumption of the services.
17 For this purpose, a provider would calculate for a given risk level ρ and
18 Quality of Service class L , the maximal number of contracts $y_{\rho,L}$ he can
19 accept. This yields the profit $e = p y_{\rho,L} - p_c C$, so that the zero-profit price is
20 $p = p_c C / y_{\rho,L}$. The provider is able to offer QoS types profitable if the
21 zero-profit price is equal or lower than market prices. Note that the
22 reverse does not hold, because multiplexing different QoS classes
23 increases often the network utilization and, in turn, reduces the costs.

1 Multiplexing gains (among different QoS types) result in portfolios with
2 $e(y) > 0$. In case that some zero-profit prices are above the market
3 prices, a portfolio y can be considered if the amount $e(y)$ can be used to
4 reduce the prices of contracts that have zero-profit prices above market
5 prices. If a provider calculates the efficient frontier, he would usually
6 eliminate the portfolios from the frontier where he would expect $e(y) <$
7 0. The reason is that in case $e(y) < 0$, the network exhibits negative
8 multiplexing gains (i.e., the assuming usage pattern cannot be allocated
9 efficiently), the network is not well suited for offering such
10 combinations of QoS classes and, hence, such combinations should not
11 be offered. A prospective provider might calculate the zero-profit prices
12 (i.e., the prices that cover costs) and the resulting base-line efficient
13 frontier. He could then compare these zero-profit prices of SLAs
14 belonging to efficient portfolios with the market prices: if all zero-profit
15 prices associated with each portfolio are, for instance, above the market
16 prices, the provider is not competitive. For a particular portfolio
17 (assuming no other financial subsidies), the losses of due contracts with
18 zero-profit prices that are higher than market prices have to be
19 compensated by profits due to contracts with lower zero-profit prices
20 than market prices.

21 Derive strategies to move towards a more efficient portfolio

22 Referring again to FIGS. 4 and 7, in order to achieve a more
23 efficient portfolio (depicted by the arrow pointing from the current

1 portfolio P^* to the efficient frontier), several options 80 for corrective
2 action are possible. In corrective action 84, a service provider might
3 reduce costs or increase risk. For this purpose, the degree of
4 multiplexing could be increased or the network capacity C decreased.
5 Note that it is sometimes possible to increase the multiplexing without
6 modifying the risk. Such a method is described by Kurz, Thiran, and
7 LeBoudec, in an article entitled *Regulation of a connection admission*
8 *control algorithm* in the Proceedings of INFOCOM'99. In corrective
9 action 82, a provider might adopt a marketing strategy to move towards a
10 more efficient portfolio. For instance, the price of the low-quality
11 service could be reduced to increase the number of contracts in this
12 class. In corrective action 92, providers might trade risks (analog to load
13 securitization and syndication): a provider can buy and sell contracts to
14 optimize his portfolio assuming that there exists a market for trading
15 contracts. For trading risks, the operator determines the number of
16 to-be-traded contracts of class i , $\Delta y_i = y_i - y_i^*$, where y_i^* and y_i denote the
17 number of contracts of class i in case of the current portfolio and a
18 desirable (i.e., efficient) portfolio, respectively. If $\Delta y_i > 0$, it's necessary
19 buy Δy_i contracts of class i , and if $\Delta y_i < 0$ the provider sells this number
20 of contracts of class i . Note that trading is a corrective action that leads
21 to an efficient portfolio (assuming that the necessary trades can be
22 executed, i.e., that there is adequate supply of SLAs having the
23 appropriate characteristics and a means of purchasing these SLAs).

24 An advantage of the invention is that it automatically and

1 rapidly calculates risk and estimated performance in transactions
2 involving network service level agreements.

3 Another advantage of the invention is that the consumer
4 may be offered a wider variety of services at a reduced price, due to the
5 associated reduction of risk brought about by better understanding of
6 risk levels for each class of services offered.

7 A latitude of modification, change, and substitution is
8 intended in the foregoing disclosure and in some instances, some
9 features of the invention will be employed without a corresponding use
10 of the other features. Accordingly, it is appropriate that the appended
11 claims be construed broadly and in a manner consistent with the scope of
12 the invention.

1 WHAT IS CLAIMED IS:

2 1. A method for managing operational risk and return of a production
3 infrastructure with respect to a current portfolio of service-level
4 agreements, the method comprising:

5 a. calculating an efficient frontier that identifies efficient
6 portfolios of SLAs using inputs such as characteristics of
7 the production infrastructure, traffic and QoS characteristics
8 and the price of each class of SLAs;

9 b. optionally, calculating a baseline efficient frontier using
10 inputs such as market pricing and break-even pricing;

11 c. determining the performance of the current portfolio of
12 SLAs using a portfolio evaluator means and inputs which
13 characterize the current portfolio; and

14 d. evaluating performance by comparing the current portfolio
15 and the efficient portfolios with the desired level of risk and
16 return; and, if desired, implementing corrective action based
17 on any desired risk and return.

18 2. The method of claim 1, wherein the corrective action is selected
19 from a group of possible actions consisting of:

20 a. adjusting marketing strategy;

21 b. changing the degree of multiplexing in the network;

22 c. changing network capacity;

- 1 d. changing the cost of network capacity;
- 2 e. defining relative compliance guarantees where networks
- 3 support definition of adequate policies on the basis of
- 4 priority;
- 5 f. changing prices and comparing with baseline prices of
- 6 SLAs; and
- 7 g. trading contracts of different classes of SLAs.

8 3. The method of claim 1 or claim 2 wherein, after corrective action is
9 taken, the method takes new inputs, and, with the exception of the
10 corrective action of trading SLAs, the method is re-executed, by
11 calculating a new efficient frontier which is compared with
12 performance of the current portfolio, calculated by the portfolio
13 evaluator means.

14

15 4. The method of claim 2 wherein, for trading risk, the operator
16 determines the number of to-be-traded SLAs of a certain class by
17 subtracting the number of SLAs of the certain class in the current
18 portfolio from the number of SLAs in a desired portfolio, and taking
19 action that tends to narrow the difference; thus moving the contents
20 of the current portfolio to that of an optimal portfolio.

21 5. A method for managing operational risk and return with respect to
22 a portfolio of service-level agreements, wherein the method uses a

1 noncompliance risk measure to calculate risk; and wherein, principals
2 of portfolio theory are applied to characterize the portfolio for
3 comparison with other possible portfolios.

4

5 6. The method of claim 5, wherein the risk measure is selected from a
6 group of quasi-linear noncompliance risk measures, the group
7 consisting of a probability of noncompliance with loss guarantees, a
8 probability of noncompliance with delay guarantees, an expected
9 penalty for loss, and an expected penalty for delay.

10

11 7. The method of claim 5 wherein the risk measure is quasi-linear and
12 the principals of portfolio theory are applied to calculate an efficient
13 frontier, thus enabling a provider to select an efficient portfolio that
maximizes return for a given risk or minimizes risk for a given return.

14

15 8. The method of claim 5, wherein the risk function is given by a
16 probability of noncompliance with loss guarantees, PNL , which, once
17 the distribution of Y , a common random variable, which represents
18 service times for customers of all classes, is known such as through
historical data, the method computes from the formula: $PNL(c, L) =$
19 $P[(Y-c)^+ - LY > 0]$, where c is C/\underline{y} , \underline{y} is the summation of a total
20 amount of accepted bandwidths of Quality of Service class L_i , C is
21 overall capacity of the network, \underline{L} is a vector which characterizes the
22 quality of each SLA, and $P[\underline{x}]$ denotes the probability of \underline{x} .

1 9. The method of claim 5, wherein the risk function is given by an
2 expected penalty for loss, EPL , which the method computes over a
3 time interval from the formula: $EPL(c, L) = (\beta C)\{E[(Y-c)+]-$
4 $LE[Y]\}$, where c is C/\underline{y} , \underline{y} is the summation of the total amount of
5 accepted bandwidths of Quality of Service class L_i , C is overall
6 capacity of the network, L is a vector which characterizes the quality
7 of each SLA, β is a constant >0 , so that (βC) denotes the penalty per
8 capacity unit, E is statistical expectation, and L_i is a total of Quality of
9 Service offered by class i .

10

11 10. The method of claim 5, wherein the risk function is given by an
12 expected penalty for delay, EPD , which the method computes over a
13 time interval from the formula: $EPD(c, L) = \beta\{(\alpha/(c-1)) - (D/c)\}$,
14 where β is a constant > 0 , $c = 1/\sum(\lambda/\mu)$, $D = c \sum\{(\lambda/\mu) D_i\}$, and
15 $E[W_i]$ denotes the expected waiting time (i.e., delay) for class i ,
16 wherein further, assumptions are made that class i traffic arrives at
17 Poisson rate λ , and that arrival processes are independent of each
18 other; service times, characterized by service rate μ of class i , are
19 independently distributed and independent of each other and of the
20 arrival processes; that $\alpha = (1 + \{Var[Y]/\mu^2\})/2$ given that service
21 times for customers of all classes are distributed as a random variable
22 Y of mean μ where $Var[Y]$ denotes the variance of random variable
23 Y , and wherein noncompliance is a penalty for exceeding D_i and a

1 premium for remaining under D_i .

2 11. The method of claim 5, wherein the risk function is given by an
3 expected penalty for delay, EPD , which the method computes,
4 assuming Poisson traffic, from the formula: $EPD(\mathbf{v}) = \beta \sum_i (E[W_i] -$
5 $D_i)$, where \mathbf{v} is a vector of traffic intensities, v_i is the traffic intensity
6 of customers in class i , E is statistical expectation, β is a constant > 0
7 so that βC denotes the penalty per capacity unit, W_i is waiting time
8 for a class i , and D_i is the maximum permissible delay for a class i of
9 SLAs.

10 12. A method for determining risk and return of a production
11 infrastructure with respect to a current portfolio, the method
12 calculating a selected risk, such as a financial risk or Quality of
13 Service risk and comprising:

- 14 a. invoking a performance evaluator means, to determine an
15 expected actual Quality of Service provided by a network
16 given a set of contracts with associated traffic descriptors;
- 17 b. calculating portfolio risk, based on the actual Quality of
18 Service and the contracted Quality of Service of the
19 contracts of the portfolio using a risk measure
20 corresponding to the selected risk; and
- 21 c. computing return according to the formula $p_i y_i - p_C C$ for
22 capacity C , expected revenue p_i , amount of contracts of

1 type i , y_i , and unit price for capacity C , p_C , where C is both
2 an input in the performance evaluator and a characteristic of
3 the production infrastructure.

4 13. The method of claim 12 wherein the performance evaluator means
5 is selected from a group of performance evaluator means consisting
6 of a formula, a simulator or test system, and a measurement system
7 for the production system.

8 14. A computerized system encoded with a method having an
9 associated process flow, the method managing the risk of financial
10 loss due to penalties brought on by noncompliance with respect to
11 network service-level agreements, characterized in that the method
12 executes the following steps:

- 13 a. gathering information such as traffic statistics, price
14 information, and network information;
- 15 b. inputting the gathered information into a risk and a return
16 function, yielding risk and return;
- 17 c. calculating an efficient frontier; and
- 18 d. using the efficient frontiers to identify an optimum portfolio
19 of service level agreements, based on a maximum level of
20 return for a given risk or a minimum risk for a given level of
21 return.

1 15. The system of claim 14, wherein, in the method, after an optimum
2 portfolio is identified, trading service-level agreements in order to
3 arrive at an optimum portfolio, the number of agreements of a certain
4 class to be traded being determined by subtracting the number of
5 SLAs of the certain class in the current portfolio from the number of
6 SLAs in a desired portfolio, and taking action that tends to narrow the
7 difference, thus moving the contents of the current portfolio to that of
8 an optimal portfolio

9

10 16. The system of claim 14, wherein, in the method, the risk function
11 is given by a probability of noncompliance with loss guarantees,
12 *PNL*, which, once the distribution of *Y*, a common random variable
13 which represents the service times for customers of all classes is
14 known such as through historical data, the method computes from the
15 formula: $PNL(c, L) = P[(Y-c)^+ - LY > 0]$, where *c* is C/y , *y* is the
16 summation of a total amount of accepted bandwidths of Quality of
17 Service class *L*, *C* is overall capacity of the network, *L* is a vector
18 which characterizes the quality of each SLA, and $P[x]$ denotes the
19 probability of *x*.

20 17. The system of claim 14, wherein, in the method, the risk function
21 is given by an expected penalty for loss, *EPL*, which the method
22 computes over a time interval from the formula: $EPL (c, L) =$
23 $(\beta C)\{E[(Y-c)^+] - LE[Y]\}$, where *c* is C/y , *y* is a summation of a total

1 amount of accepted bandwidths of Quality of Service class L_i , C is
2 overall capacity of the network, \underline{L} is a vector which characterizes the
3 quality of each SLA, β is a constant >0 , so that (βC) denotes the
4 penalty per capacity unit, E is statistical expectation, and L_i is the
5 total Quality of Service offered by class i.

6

7 18. The system of claim 14, wherein, in the method, the risk function
8 is given by an expected penalty for delay, EPD , which the method
9 computes over a time interval from the formula: $EPD(c, L) = \beta \{ (\alpha$
10 $/(c-1)) - (D/c) \}$, where β is a constant >0 , $c = 1/\sum(\lambda/\mu)$, $D = c \sum$
11 $\{(\lambda/\mu)D_i\}$, and $E[W_i]$ denotes the expected waiting time (i.e., delay)
12 for class i , wherein further, assumptions are made that class i traffic
13 arrives at Poisson rate λ_i , and that arrival processes are independent
14 of each other; service times, characterized by service rate μ of class i ,
15 are independently distributed and independent of each other and of
16 the arrival processes; that $\alpha = (1 + \{Var[Y]/\mu^2\})/2$ given that service
17 times for customers of all classes are distributed as a random variable
18 Y of mean μ where $Var[Y]$ denotes the variance of random variable
19 Y , and wherein noncompliance is a penalty for exceeding D_i and a
20 premium for remaining under D_i .

21 19. The system of claim 14, wherein, in the method, the risk function
22 is given by an expected penalty for delay, EPD , which the method
23 computes, assuming Poisson traffic, from the formula: $EPD(v) =$

1 $\beta \sum_i (E[W_i] - D_i)$, where v is a vector of traffic intensities, v_i is the
2 traffic intensity of customers in class i , E is statistical expectation, β
3 is a constant > 0 so that βC denotes the penalty per capacity unit, W_i
4 is waiting time for a class i , and D_i is the maximum permissible delay
5 for a class i of SLAs.

6 20. A computerized system encoded with a method which manages
7 operational risk and return with respect to network service-level
8 agreements, wherein the method calculates a probability of actual loss
9 higher than allowed by a contract and a return, and, applying the
10 principals of portfolio theory, determines an efficient frontier to
11 enable the selection of an efficient portfolio that maximizes return at
12 a given risk or minimizes risk at a given return.

13 21. The system of claim 20 wherein, in the method, the return is
14 calculated using an expected penalty for loss.

15 22. A computerized system, encoded with a method executing a
16 process flow which manages operational risk and return with respect
17 to network service-level agreements, operating over a computer
18 network comprising a plurality of interconnected computers and a
19 plurality of resources, each computer including a processor, memory
20 and input/output devices, each resource operatively coupled to at
21 least one of the computers and executing at least one of the activities

1 in the process flow, wherein the method manages a portfolio of
2 service level agreements, each of which define a service level, a
3 connection, a contract duration, traffic descriptors, quality of service
4 guarantees and a probability of noncompliance with respect to the
5 quality of service guarantees, the probability of noncompliance
6 providing a contractual parameter wherein, after being accepted by a
7 customer, noncompliance within the contracted limits does not trigger
8 a penalty, thus avoiding penalties for noncompliance and thus
9 reducing.

10 23. The system of claim 22, wherein, the quality of service guarantees
11 include loss rate, delay, and jitter.

12 24. A computerized system encoded with a method which manages
13 operational risk and return with respect to service-level agreements in
14 a network, wherein the method manages a portfolio of service level
15 agreements of at least two classes each of which representing relative
16 compliance guarantees, wherein, a customer subscribing to a higher
17 relative compliance guarantee has priority with respect to resources in
18 the network, over customers having a lower relative compliance
19 guarantee.

20 25. A computerized system encoded with a method which manages

1 operational risk and return with respect to network service-level
2 agreements, wherein the method takes probabilities of noncompliance
3 and base-line prices, and, through the application of portfolio theory,
4 calculates an efficient portfolio of service-level agreements, thus
5 providing a network administrator with insights into the economics of
6 a network's operations which can be used to modify the terms of
7 standard service-level agreements.

8 26. The system of claim 25, wherein the base-line prices are
9 zero-profit prices.

10 27. The system of claim 25, wherein the base-line prices are market
11 prices.

12 28. The system of claim 26, wherein the zero-profit prices are
13 calculated by:

- 14 a. calculating a base-line efficient portfolio using market
15 pricing and thus determining base-line prices;
- 16 b. investigating which of these portfolios are probably
17 attainable;
- 18 c. comparing the base-line prices against a zero-profit price;
- 19 d. if the zero-profit price is higher than the base-line price,
20 taking corrective action.

21

1 Portfolio Theory Method of Managing Operational Risk
2 with Respect to Network Service-level Agreements

3 Abstract of the Invention

4 A method for managing operational risk and return with respect to a
5 portfolio of service-level agreements is provided, wherein the method
6 uses a noncompliance risk measure to calculate risk; and wherein,
7 principals of portfolio theory are applied to characterize the portfolio for
8 comparison with other possible portfolios.

U.S. GOVERNMENT USE

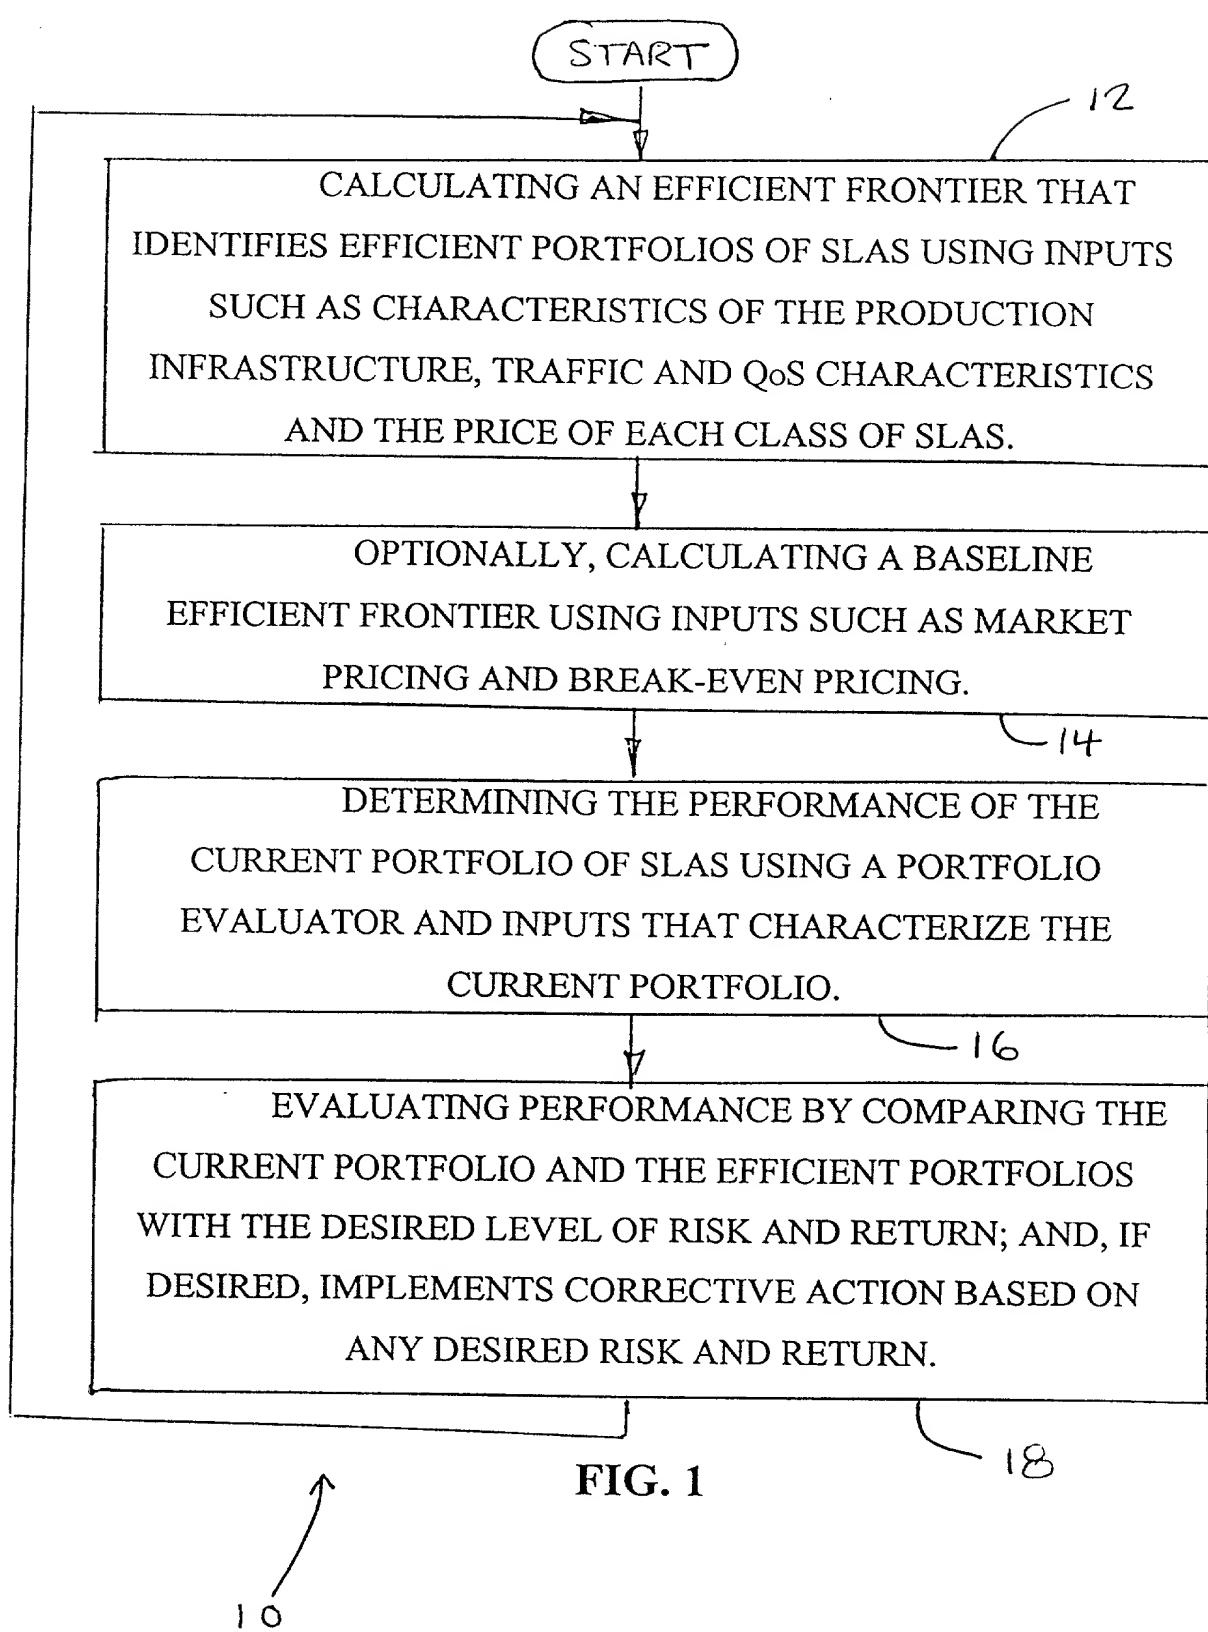


FIG. 1

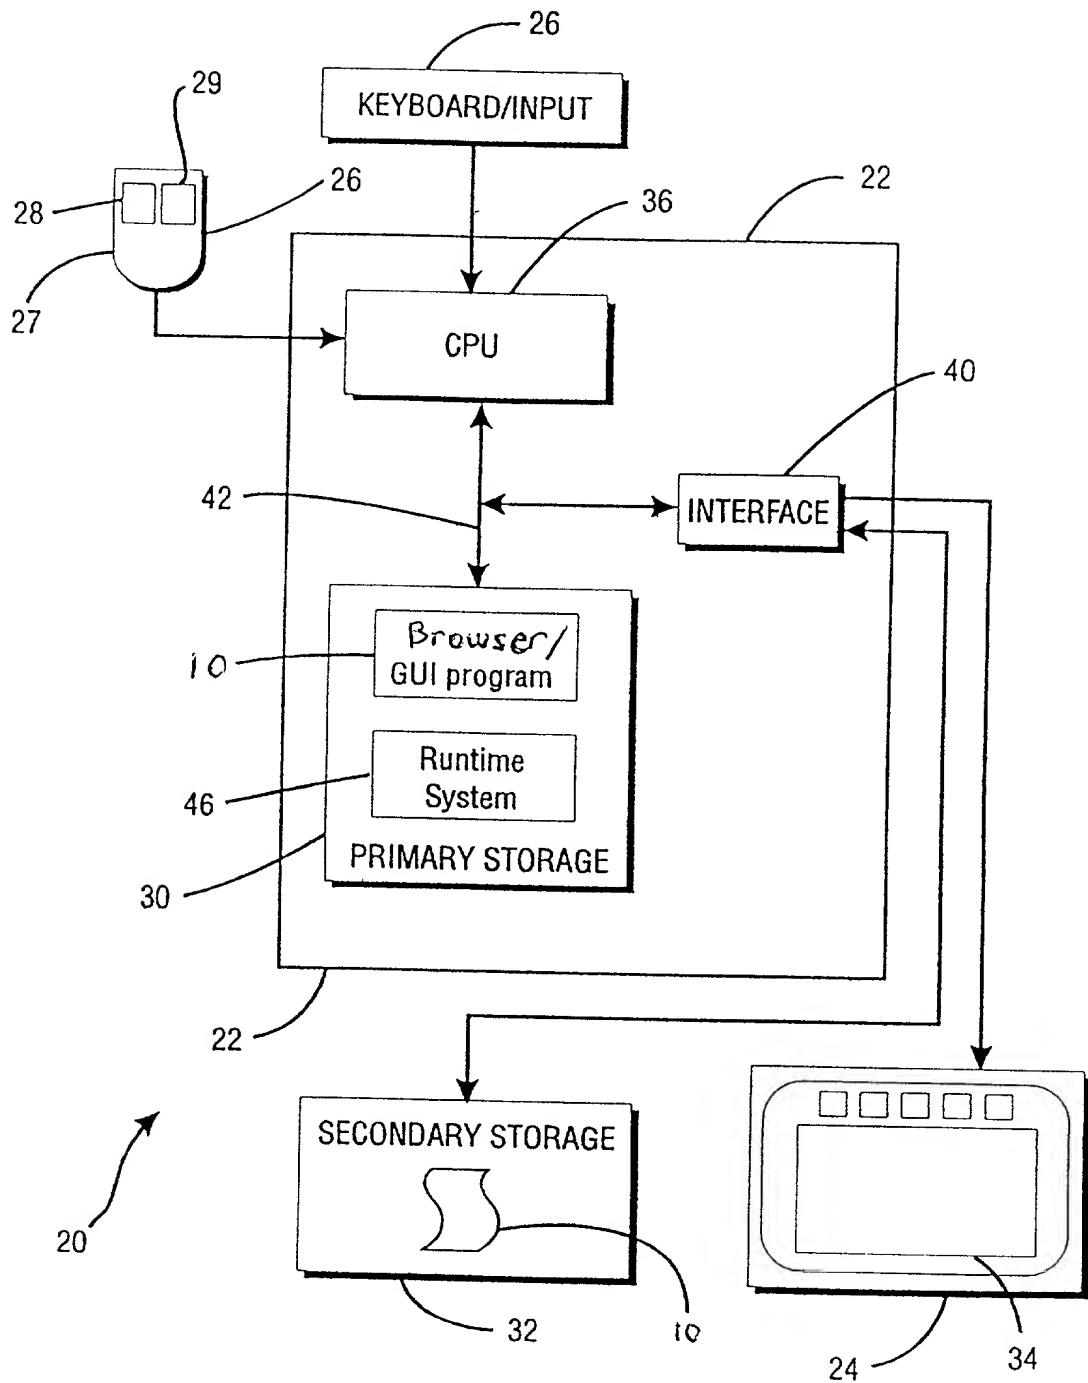


FIG. 2

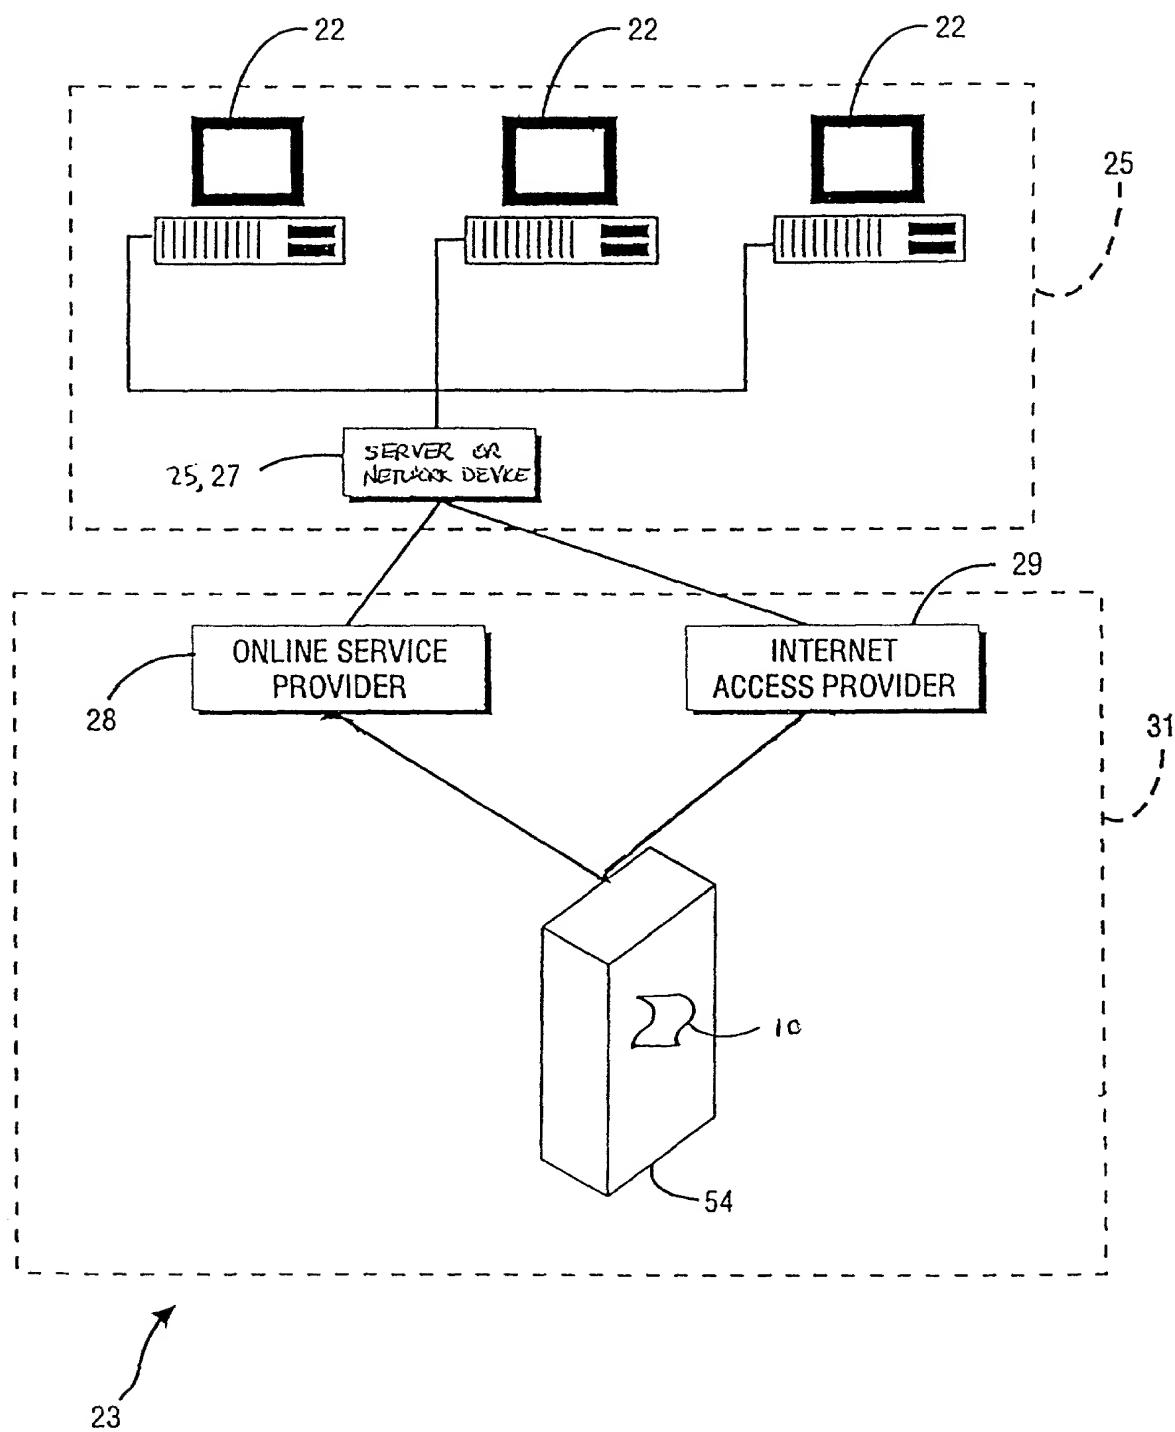


FIG. 3

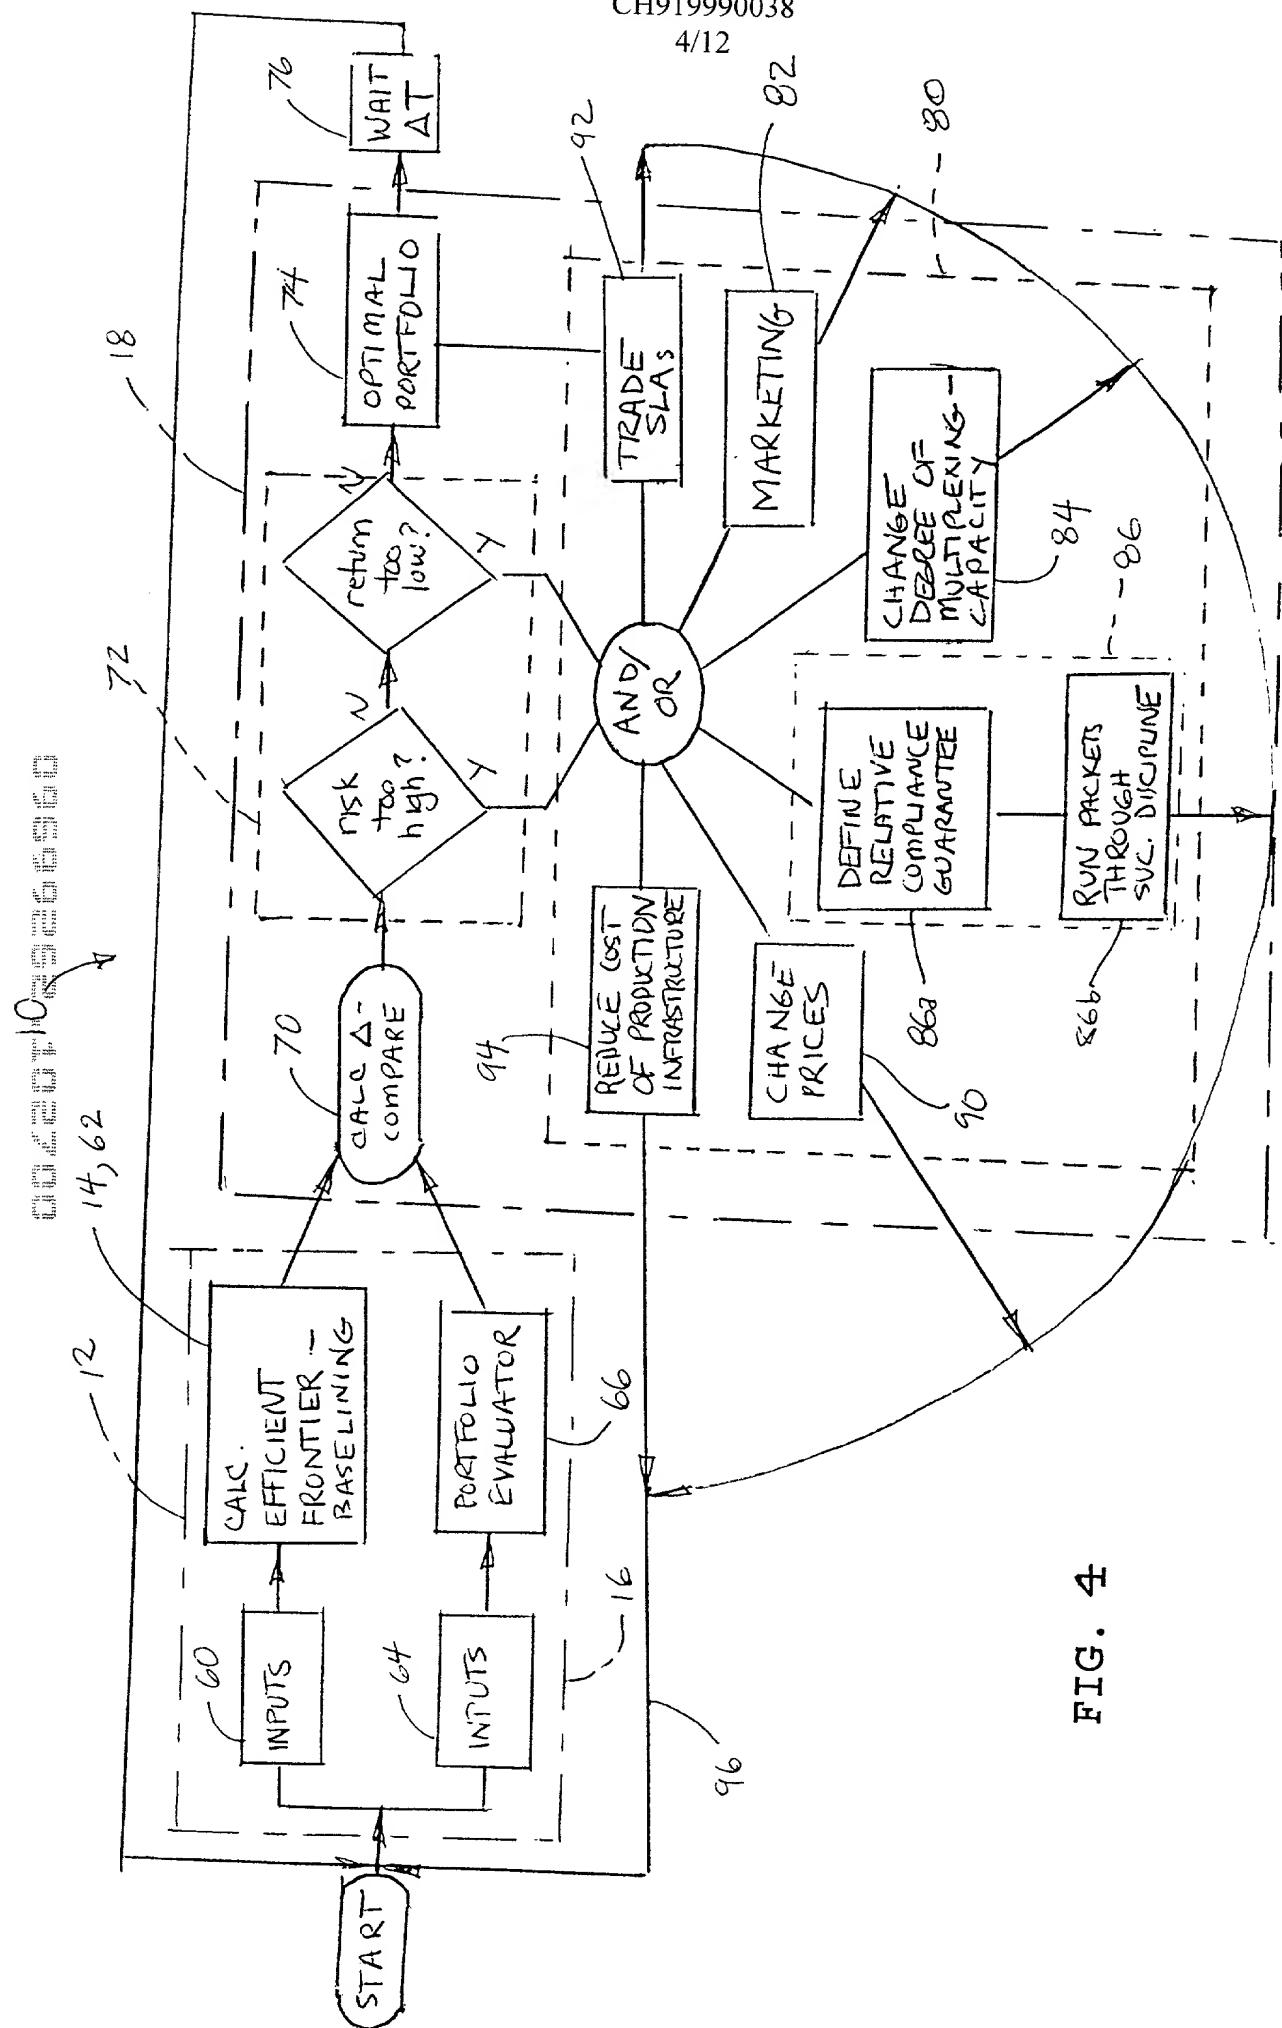


FIG. 4

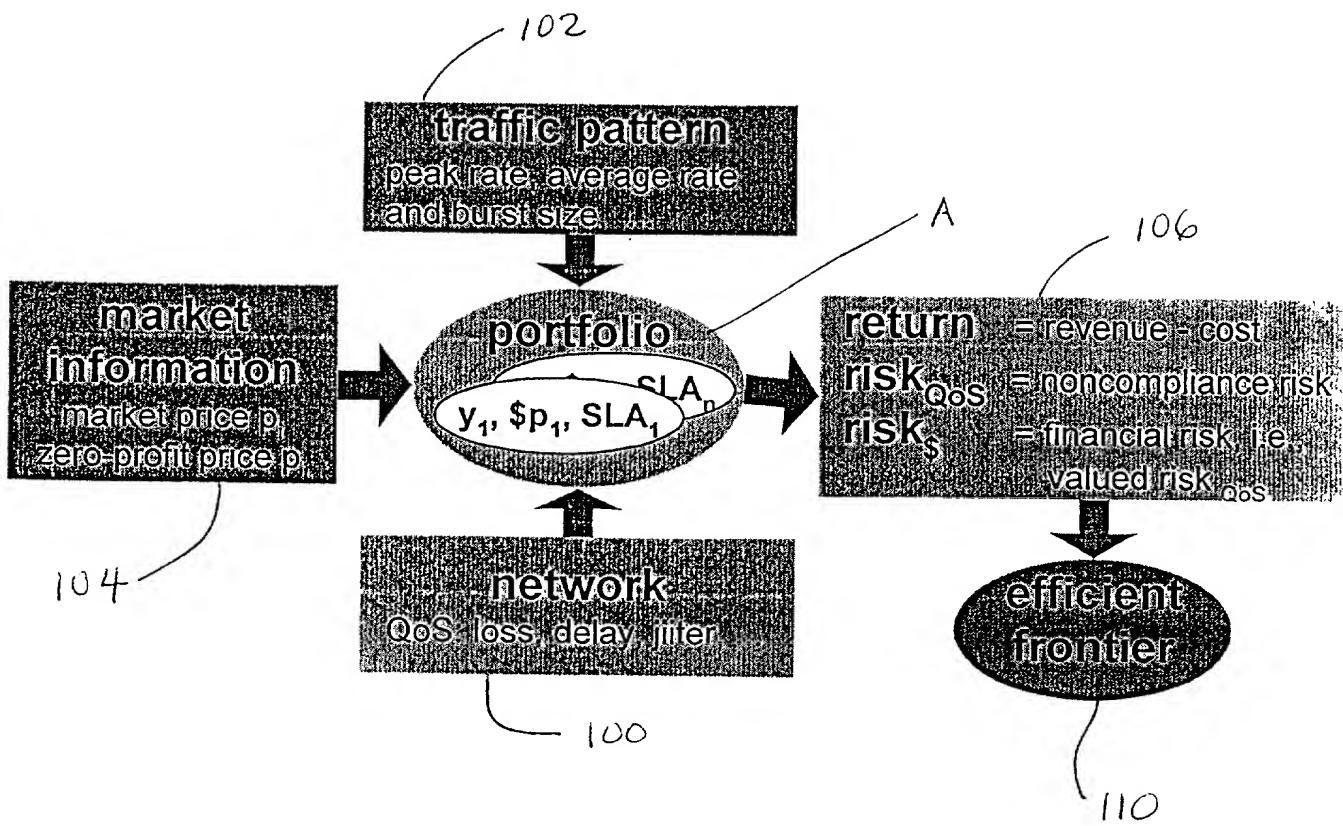


FIG. 5

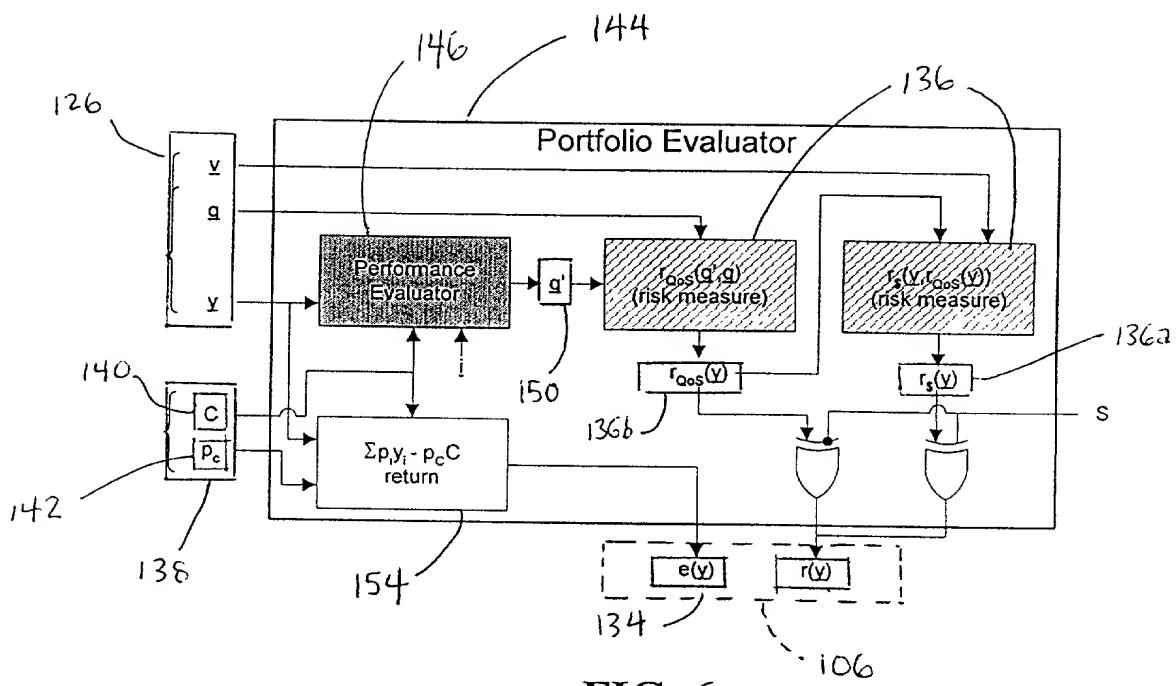


FIG. 6

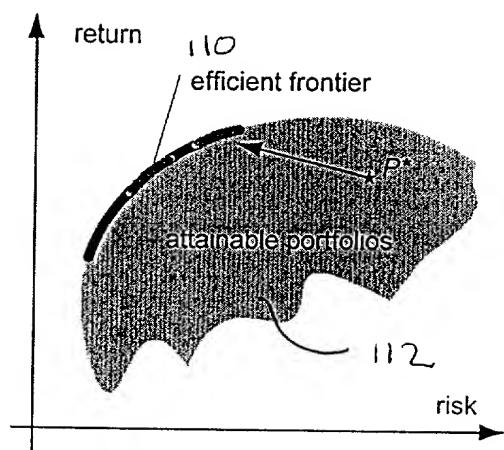


FIG. 7

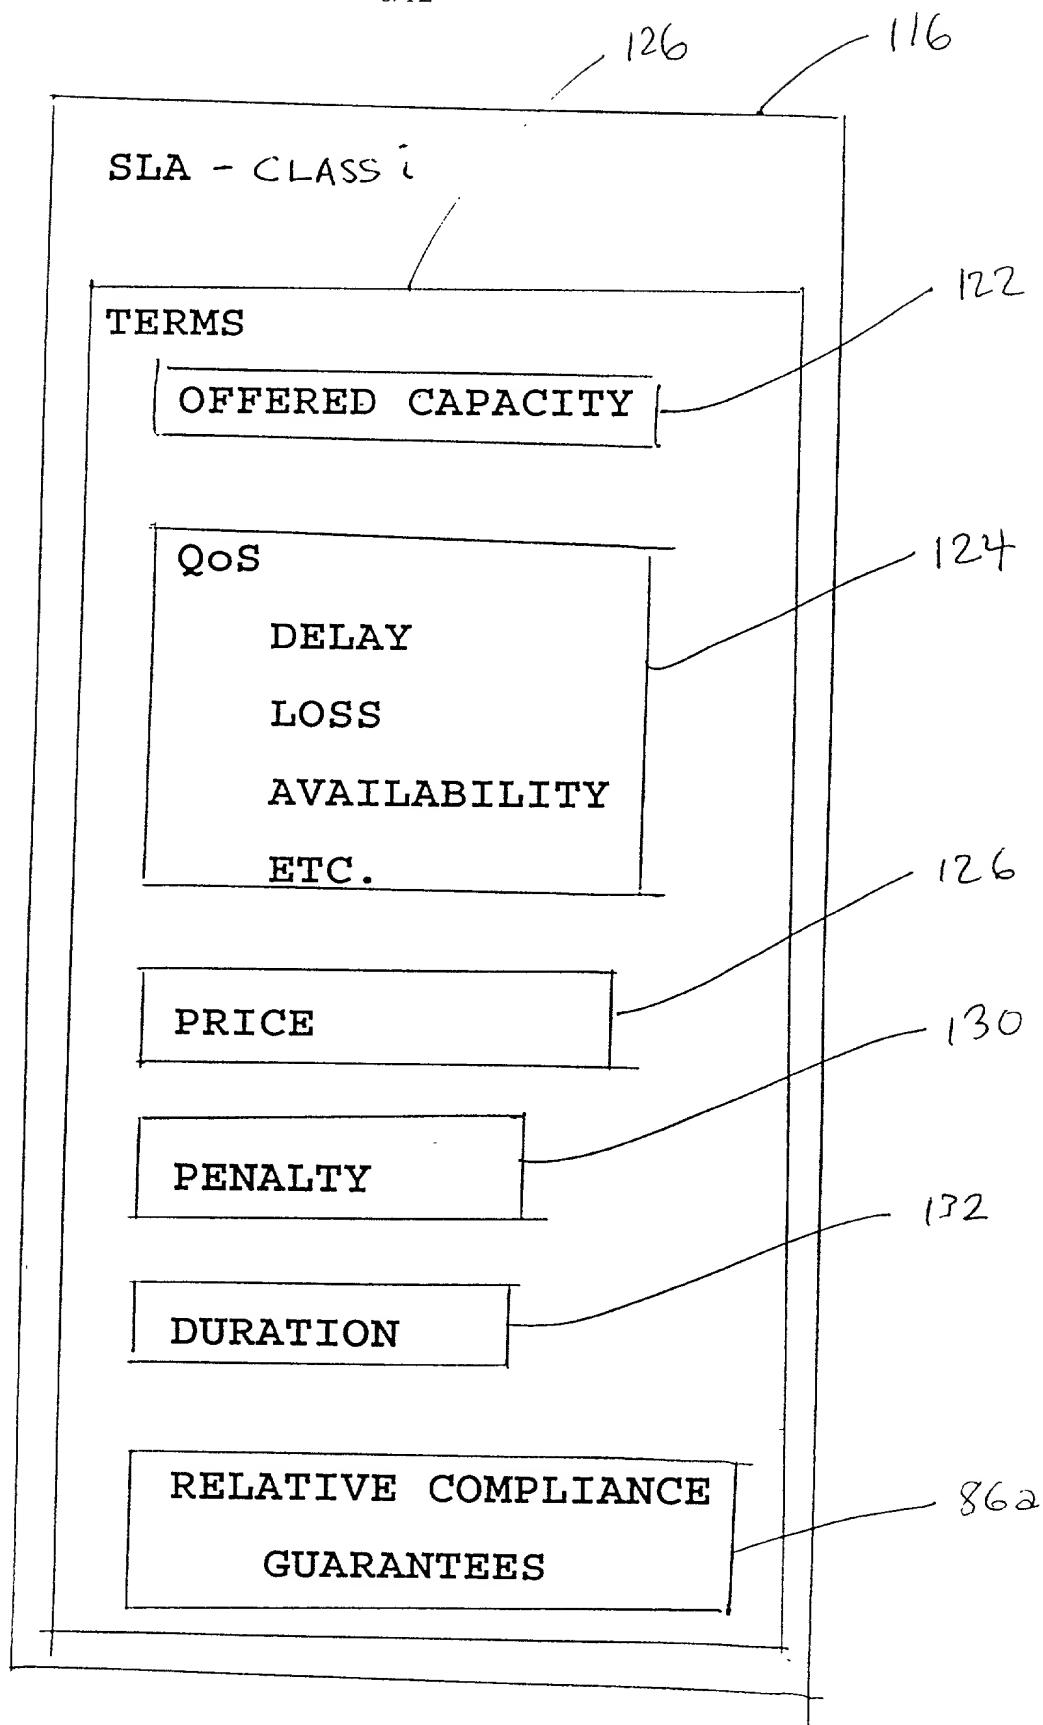


FIG. 8

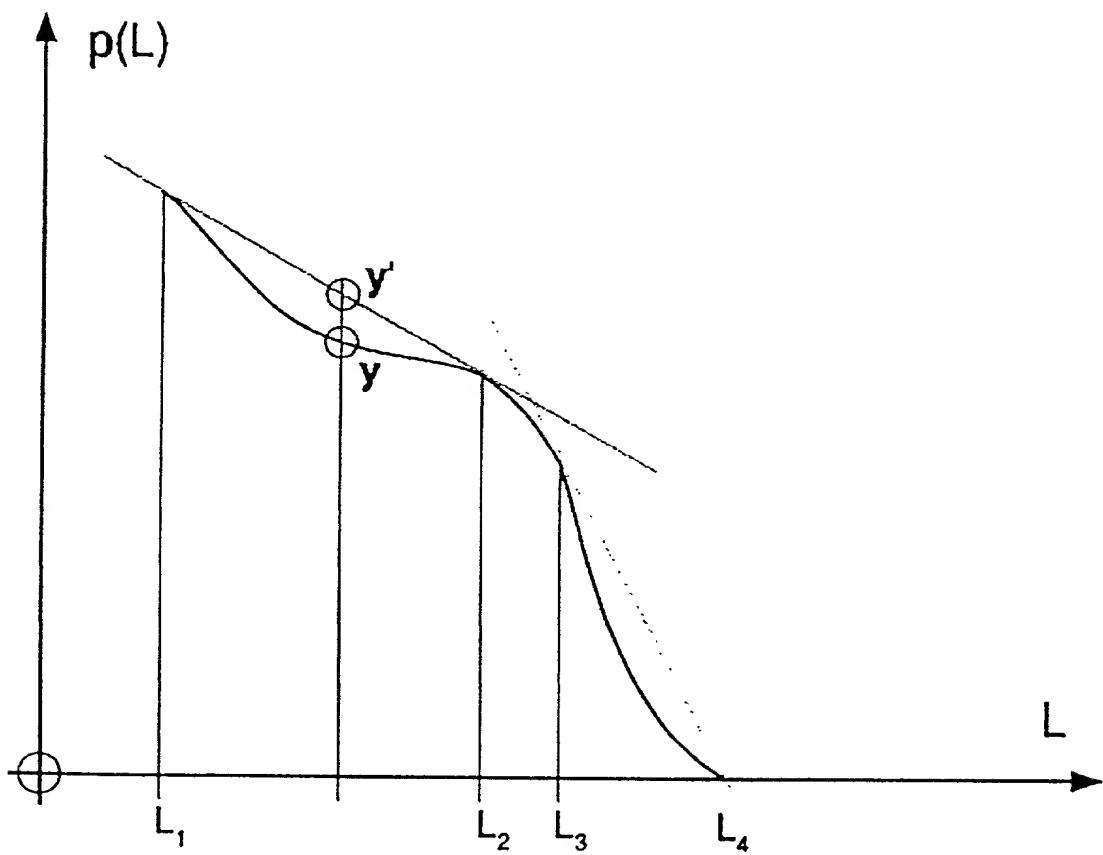


FIG. 9

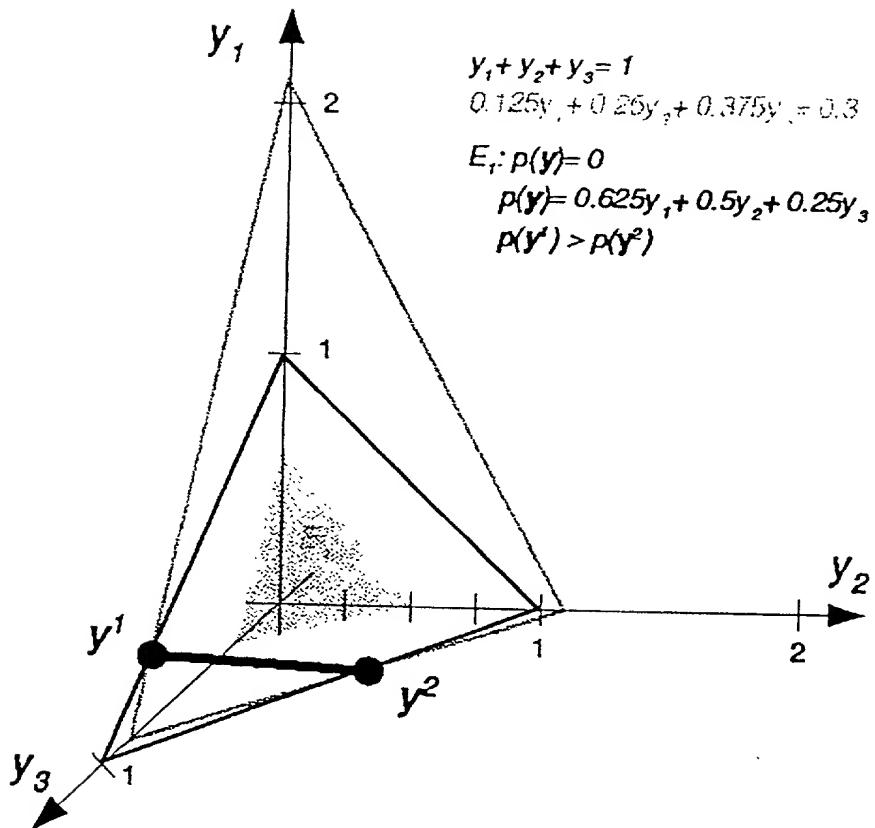


FIG. 10

CH919990038

11/12

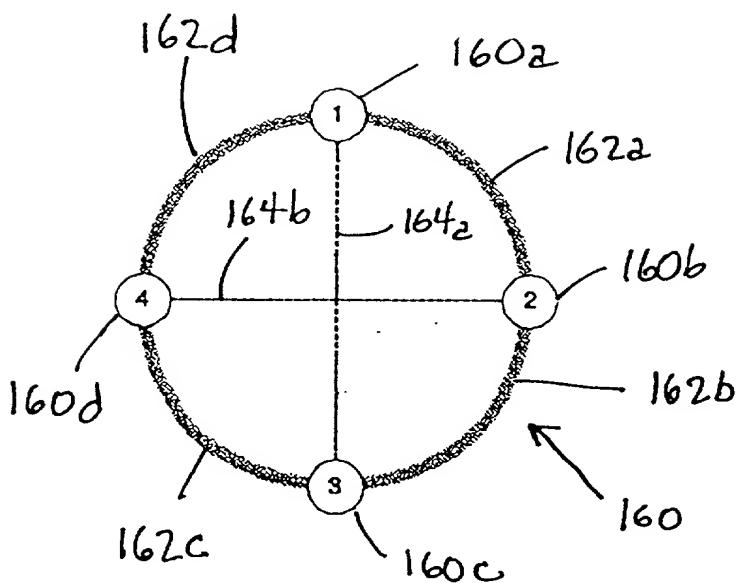


FIG. 11

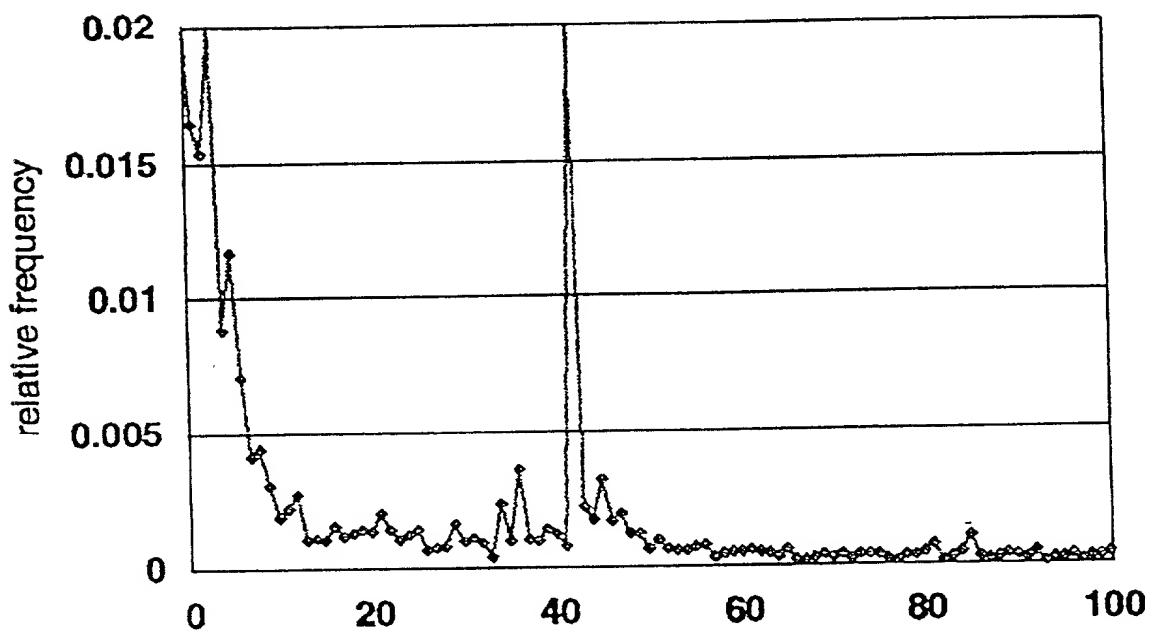


FIG. 12

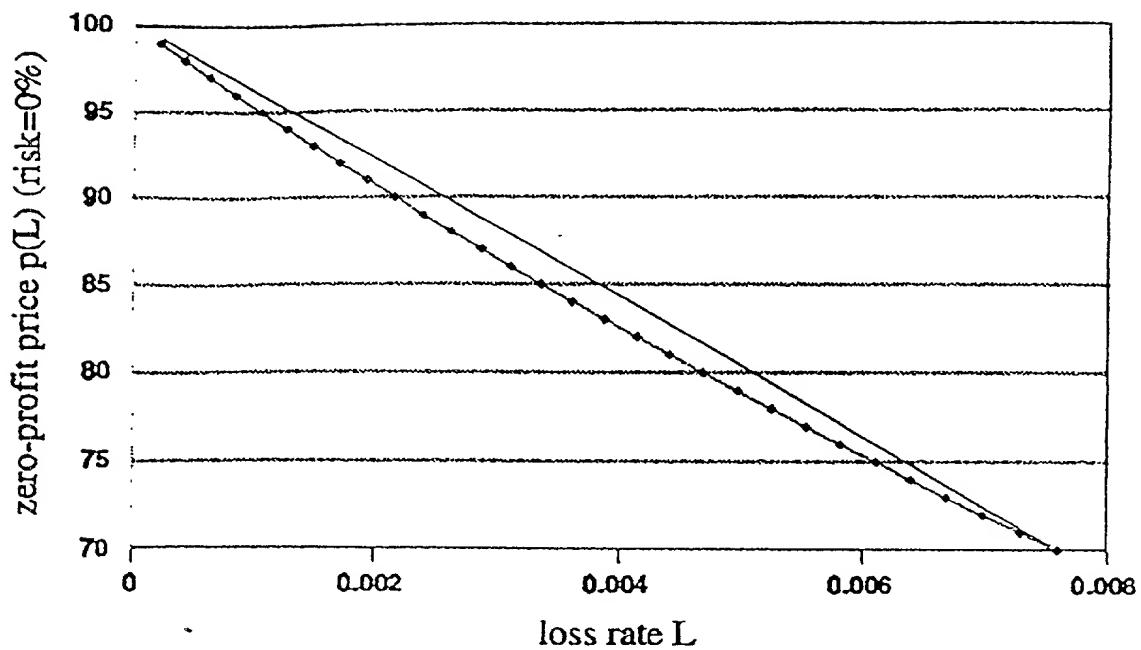


FIG. 13

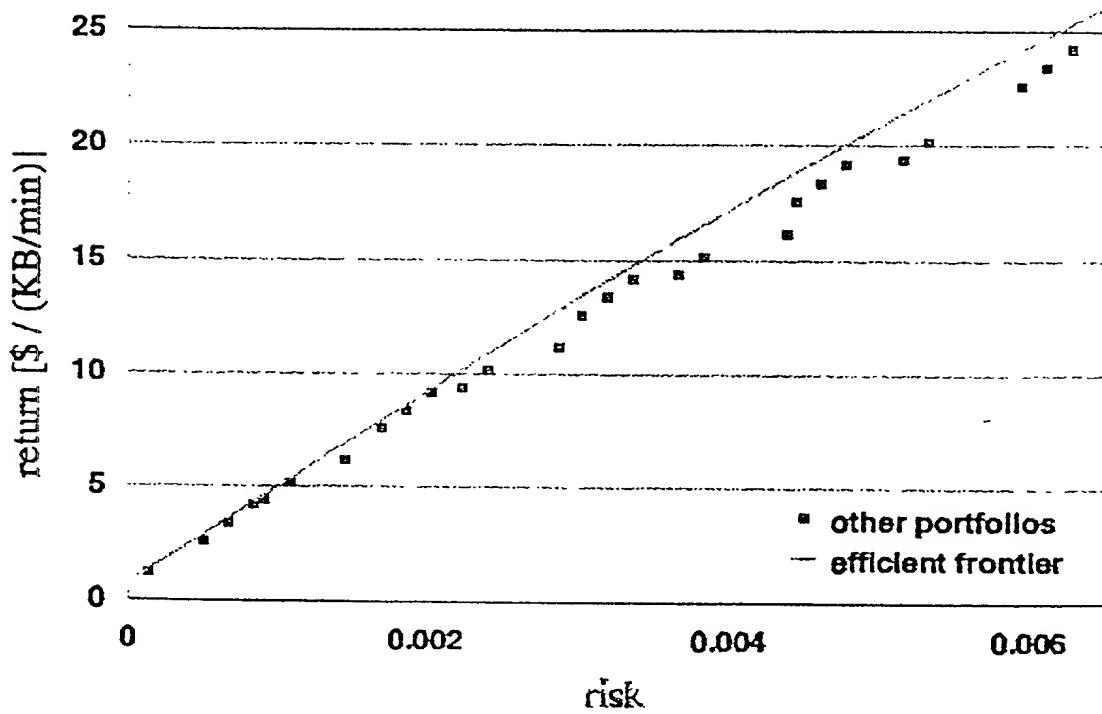


FIG. 14

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATIONExpress Mail Label EL549238338US
Date of Deposit October 27, 2000Beat LiverFull name of second joint-inventor, if any

Inventor's signature

Date

Ruetistrasse 36, CH-8134 Adiswil, Switzerland
ResidenceSwiss
CitizenshipSame as above.
Post Office AddressFull name of third joint-inventor, if any

Inventor's signature

Date

Residence

Citizenship

Same as above.
Post Office AddressFull name of fourth joint-inventor, if any

Inventor's Signature

Date

Residence

Citizenship

Same as above.
Post Office AddressFull name of fifth joint inventor, if any

Inventor's Signature

Date

Residence

Citizenship

Same as above.
Post Office AddressFull name of sixth joint-inventor, if any

Inventor's signature

Date

Residence

Citizenship

Same as above.
Post Office Address